Innovative use of Data Networks to Improve Campus Security

JTAP-633

Mark Toller
Tim Chown
Paul Lewis
John Clark

Department of Electronics and Computer Science
University of Southampton

January 2000

1 Contents

2 Executive Summary

This report considers three main aspects of security. It looks at legal and social aspects of deploying a CCTV security system. It investigates a number of commercial systems for CCTV over a variety of media including data networks, and explores the deployment and effectiveness of such systems. Thirdly, avenues for a more "complete" security system are explored.

The primary conclusions of the work, carried out over the period of a year during 1998/9 are:

1. The installation of a successful security system involves not only the deployment of a CCTV system but also a wider more "holistic" view of security issues, including physical alarms, secure access mechanisms (such as card lock devices) and staff awareness. By adopting such a view, reported incidents of theft at the authors' site have been significantly reduced.

2. Category 3, 5 or 5+ twisted pair cabling can be used for a variety of purposes. Traditionally it has been used for telephony (including ISDN), serial connections and data, but it can also be used for alarm wiring and raw video transmission. When designing a new building, specifying extra twisted pair outlets for alarm wiring can offer a flexible and cost effective basis for a reliable security system. In older buildings deploying Category 5+ cable, the older existing Category 3 cable can be re-used for alarm wiring.

3. A low-cost CCTV system can be deployed in a typical Departmental scenario. However, given the lack of effective Quality of Service (QoS) mechanisms on Ethernet networks, it is still most likely to be more cost-effective to operate a high quality CCTV system over dedicated cable (coaxial, twisted pair or, if distance requires it, fibre-optic).

4. The low cost of well-specified PCs, digitising cards and, in particular, large IDE hard disks makes "digital" recording of images an attractive proposition. With just the addition of a cheap (around £70) parallel port camera, any PC can become a WebCam triggerable by movement (image differencing), capable of uploading images to an FTP server or saving (lower quality) video sequences to disk. However, correlation of images or video from multiple sources into a single queryable database is a non-trivial task.

5. CCTV and alarm systems can be used in areas beyond deterrence of theft. With many staff and students undertaking lone working in offices or laboratories, CCTV may offer a method to help monitor them, and to offer additional safety in situations such as electronics laboratory sessions. The same cameras can be used for security purposes by night.

6. The 1998 Data Protection Act has new implications, even for those sites deploying small-scale or single-camera surveillance systems. In contrast to the 1984 DPA, data subjects have improved rights, such that almost any CCTV system is now liable for registration under the new Act. Computer system managers should be aware that they are also liable for offences under the 1998 DPA where the CCTV footage is handled digitally and/or over data networks.

7. Digital images and video can be used in evidence in a court of law. Any image may be used in evidence, but the weight it will be given depends on a number of factors, not least the establishment of an audit trail and also the authenticity of the image(s).

Recommendations and suggestions for further work can be found at the end of the report.

3 Introduction

This report is based on work undertaken within the Department of Electronics and Computer Science (hereafter "the Department") at the University of Southampton. While the scope of the project included the whole Department, much of it focused on measures introduced in the New Zepler building, which was formally commissioned in the Spring of 1999.

The University of Southampton lies on the south coast of England and has some 25,000 staff and students. The Department is one of the largest Electronics and Computer Science Departments in the UK, with some 1,200 computer users including over 200 new degree students per year. The New Zepler building is a four-story building housing a mixture of academics, student laboratories and administrative staff.

There are two main aspects to this report. The first is non-technical from a computing point of view, and investigates issues surrounding the deployment of a security (CCTV in particular) system. This includes notes on the new 1998 Data Protection Act (DPA) and comments on the use of digital video evidence in a UK court of law. The second aspect describes the technical trials and deployments undertaken in the Department, including both CCTV and other security measures.

The report has a number of recommendations for HE Institutions, or Departments at such institutions, who are considering deploying security systems. It also includes suggested measures for taking a more "holistic" view of security.

The preceding Executive Summary describes our major findings. Recommendations, for Institutions and for further study, are to be found at the end of the report.

3.1 Acknowledgements

The work described in this report was made possible by twelve months of funding by the Joint Information Systems Committee (JISC) of the Higher Education Funding Councils.

Other people whose help has been valuable include but are not limited to:

Adrian Pickering, Department Manager
Shaun Ford, Department Technician
Colin Bird, IBM
Mark Hindess, University of Bath
David Owen, University of Southampton Chief Security Officer
Bob Ballard, IT Guard
Jeff Anderson, ACI International
Tom Franklin and Tish Roberts, JISC
The Hampshire Constabulary
Mark Thompson and his digital camera

4 Data Protection Act Issues for CCTV and PC Cameras

The Data Protection Act 1984 has recently been updated in the 1998 Act to conform to the EU Data Protection Directive (95/46/EC). The new Act received Royal Assent on July 16th, 1998. It will come into force on March 1st, 2000. It is important that we understand the new DPA issues regarding both analogue or digital CCTV and PC Camera, or WebCam, use.

A WebCam is a networked device (usually a PC with a digitising card) with an attached camera via which either still images or live video can be stored and/or made available to other network users, in the case of images typically on a World Wide Web site. The low cost and resulting proliferation of parallel port or USB-based PC cameras and WebCam software makes the legal and proper use of such devices a real concern for computer and systems managers.

Because the new Act is not yet in force at the time of writing, and there have thus as yet been no test cases, we can only highlight some potential areas of importance to Institutions and their data controllers.

4.1 The 1984 Data Protection Act

The Data Protection Act 1984 applies to the use of certain CCTV systems, but several did not fall within the specified rules and did not have to register. There would certainly have been very few instances of CCTV footage being used in digital form on computers in 1984, given that specialised "frame grabbers" were both expensive and not easy to operate. The 1998 Act has several changes from the 1984 DPA, many of which recognise the new electronic medium, though the majority of compliance with the 1998 Act follows from complying with the original 1984 Act.

The 1984 DPA, which requires that data be held "fairly and lawfully", asked three main questions to determine whether or not you (as the data user) needed to register your CCTV system. These were:

1. Am I collecting data?
2. Am I collecting personal data?
3. Is the personal data automatically processed by reference to the individual?

If the answer to all three questions was "yes" then your CCTV system needed to be registered.

In terms of CCTV surveillance, it is necessary to answer "yes" to question 1 if either video or image data is being recorded. Whether this data is stored on a computer or recorded on video tape is irrelevant - both are covered by the Data Protection Act 1984.

It is in questions 2 and 3 where the 1984 Act fails to cover many situations arising from advances in digital technology. For question 2 "personal data" as defined by the 1984 Act means

"data consisting of information which relates to a living individual who can be identified from that information (or from that and other information in the possession of the data user), including any expression of opinion about the individual but not any indication of the intentions of the data user in respect of that individual"

Thus, if you capture images of people unknown to you in a public place, you are not collecting "personal data" and will thus most likely not have to register your system. You will only have to register if you have another means at your disposal to identify that person (for example a correlation with card lock logs where a camera is surveying a door – i.e. you can make an identification by reference to the individual). If the camera is at your place of work, and you may know the employees passing the camera, then you are potentially collecting personal data and you may need to register.

For question 3, the 1984 Act defines processing as

"amending, augmenting, deleting or re-arranging the data or extracting the information constituting the data and, in the case of personal data, means performing any of those operations by reference to the data subject"

In order to view or print information (or any other processing) regarding an individual from CCTV data, then extraction must have taken place, and therefore the data must have been processed. However, a CCTV system will only fall within the provisions of the Act if it can automatically locate particular information stored on the system. In the guidelines this is worded as follows:

"If your system can be programmed to search for a specific time or frame reference on the tape at which it is known that personal data about a known individual is recorded, then your equipment will be capable of automatically processing."

If you are performing automatic processing, you will need to register, assuming you also answered "yes" to questions 1 and 2. Failure to register when required to do so is a criminal offence.

4.2 The 1998 Data Protection Act

In the House of Lords Select Committee on Science and Technology, Eighth Report on Digital Images as Evidence, Appendix 2, the Lord Williams of Mostyn (Parliamentary Under Secretary of State, Home Office) describes the weaknesses of the 1984 Act, and how the 1998 Act is intended to rectify it:

"*As the Data Protection Registrar explained in her oral evidence, a number of CCTV systems are outside the scope of the 1984 Act because of the limitations of its definition of personal data and its requirement for processing to be undertaken by reference to the individual. Clause 1(1) of the Bill significantly expands both definitions. It appears to us that monitoring by live CCTV comes within the first limb of the definition of "data", i.e. information which is being processed by means of equipment operating automatically in response to instructions given for that purpose. "Processing" is no longer limited to activities carried out by reference to the data subject and it also covers a much wider range of operations, including the disclosure of data by transmission, dissemination or otherwise making available. We consider that the new definitions are sufficiently broad to catch not only the sophisticated types of CCTV systems with which the Committee are concerned, but even much simpler equipment which merely projects images of individuals passing a shop into the shop window without actually recording*."

The text goes on to state:

"*You queried firstly whether the digital television surveillance systems which are increasingly likely to be used in city centres were unambiguously within the scope of the proposed legislation. We are confident that they are*."

But what is the nature of the extended scope of the Act? According the the UK Government data protection Web site:

"The new Act contains familiar elements from the current legislation such as; the Data Protection Principles of good practice; a registration system; an independent supervisory authority to oversee data protection legislation; and the data subject's right to have access his or her personal data and to correct it where inaccurate. However the [EU] Directive does impose additional requirements which are also reflected in the new law…

The 1998 Data Protection Act goes further still, and now covers not only digitally stored information, but also manually stored records (e.g. paper-based filing systems). However, the effect will not be immediate:

"*Manual data and automated data (where the 'processing was already under way immediately before 24 October 1998) will not have to comply with the new Act until 2001, and data which was already held in manual filing systems prior to the 24 October 1998 need not comply with some aspects of the new Act until 2007*."

The 1998 DPA is no longer limited to processing operations that are carried out by reference to the data subject directly. In the new Act personal data now means

"*data which relate to a living individual who can be identified from those data, a combination of data in the data controller's possession, or from information that is likely to come into the possession of the data controller, and includes any expression of opinion about the individual and any indication of the intentions of the data controller or any other person in respect of the individual*."

The definition of processing is now expanded as well:

> *""processing" in relation to information or data, means obtaining, recording or holding the information or data or carrying out any operation or set of operations on the information or data, including-*
>
> *(a) organisation, adaptation or alteration of the information or data,*
> *(b) retrieval, consultation or use of the information or data,*
> *(c) disclosure of the information or data by transmission, dissemination or otherwise making available, or*
> *(d) alignment, combination, blocking, erasure or destruction of the information or data;"*

Disclosure could thus include via Web pages or electronic mail (e-mail).

Data controllers (who determine the purposes for which and manner in which any personal data is used) continue to have an obligation to take security measures to safeguard personal data, and to take measures to prevent unauthorised processing of the data. Data subjects have enhanced rights of access, including a description of the data being processed and the purpose of the processing. They also have the right to prevent data processing that may cause "damage or distress" (and to seek compensation if it is judged that it does). "Distress" may presumably

include the display of digitised images in a public place, be that on a TV screen, a Web page or in printed form. And, as with the 1984 Act, personal data may not be kept for longer than is necessary to meet the stated purpose of holding the data.

Section 29 of the new Act has limited exemptions for "crime and taxation", but these do not prevent the CCTV video or image data from being required to be registered; they appear to be more related to how the data can be gathered (for example the use of hidden cameras).

The provisions of the new Act make clear the need for caution when considering deploying even a single WebCam device in a room, especially if the data is then placed on a Web page.

There is also one further change in the new Act that may impact the Internet rather directly, and which may reinforce the need for caution:

"In addition to the changes to the existing Principles outlined above there is also a new eighth principle restricting the transfer of personal data outside the EU. There are to be no restrictions on the free flow of personal data between countries in the European Economic Area (which consists of Norway, Iceland and Liechtenstein, as well as the 15 EU Member States). However, personal data may only be transferred to third countries if those countries ensure an "adequate level of protection for the rights and freedoms of data subjects"."

An image displayed on a Web page, viewed from the USA (which we believe has not yet been stated to have that "adequate level"), could be seen to fall into this category. Such an image is eminently downloadable and storable by the viewing party.

4.3 Liabilities

Section 61 of the 1998 DPA makes clear the liability for offences:

"*Where an offence under this Act has been committed by a body corporate and is proved to have been committed with the consent or connivance of or to be attributable to any neglect on the part of any director, manager, secretary or similar officer of the body corporate or any person who was purporting to act in any such capacity, he as well as the body corporate shall be guilty of that offence and be liable to be proceeded against and punished accordingly*."

Computing managers may thus potentially be liable in the case where images or video is held on or transmitted over a public University system, e.g. a University Web server.

4.4 Conclusion

The 1998 Data Protection Act expands the scope of what constitutes a system that must be registered, in terms of what is meant by personal data and processing, such that almost any video surveillance system will likely have to be registered. Registering a system has implications for what can be done with the data (e.g. the restriction on data transfer outside the EU).

Individuals now have improved rights, and may be able to request information that would not have had to been made available under the old Act. Institutions should consider re-evaluating their definitions of the data they hold and the processing performed (as the subject may request such information) as well as how such data as (public) CCTV images can be made available when requested by a data subject.

The full situation is currently unclear due to a lack of test cases. Given that managers are liable for offences along with their Institution (with a Level 5 fine in the Act being anything up to £5,000), they should consult their University Data Protection Officer (where present) as soon as possible.

There is a mailbase mailing list *data-protection@mailbase.ac.uk* which discusses DPA issues.

5 Digital Images as Evidence

There is little case law on the subject of digital images being used as evidence, generally the cases where computer evidence are used concern computer logs, printouts, emails and the like - i.e. text documents either stored on, or generated by a computer. As digital CCTV and other digitally captured and stored images become more prevalent, this issue is a very important one. Probably the most authoritative reports on the subject are the Fifth and Eighth Reports from the House of Lords Select Committee on Science and Technology on "Digital Images as Evidence". The Fifth Report was published on the 21st February 1998, and the Government produced a reply to it on the 22nd April 1998 from which the Eighth Report arose. That report is aimed at all aspects of digital images, i.e. document imaging, traffic control, overt and covert surveillance, but the majority is based on CCTV digital surveillance images.

5.1 Authenticity

One of the main problems with this new technology is verifying the authenticity of a digital image. Due to improvements in imaging hardware and software, it has become increasingly possible to make a modified image which "*may be indistinguishable from the first version in terms of quality or the apparent representation of reality*". This, coupled with the fact that a digital image may be copied exactly, modified, and then written to an identical storage device (such as a CD) can cause concern for the Crown Prosecution Service. The ability to make an exact copy of a digital image also causes some confusion of what can be described as the original. In the report above, the following definition is used:

"*We prefer the following definition in relation to digital images: the original is the data first recorded in memory. Thus any printed or displayed image created from these data is a copy. Consequently digital recording technology provides no original that could be produced in evidence. All that is available for use as evidence is a copy of the first, probably temporary, recording in memory, and this will be admissible as evidence. Its weight as evidence will depend on proper authentication and other matters*."

5.2 Proper operation of the computer

Under the law of England and Wales, "*if the original of a document no longer exists, copies or even copies of copies are admissible in evidence and it is irrelevant that the original was destroyed by the person seeking to produce the copy as evidence*". It would seem, therefore, that digital images can always be admitted as evidence in a court of law, but there is a caveat - in a criminal case, the provisions of Section 69 of the Police and Criminal Evidence (PACE) Act 1984 apply to all computer records:

"1) In any proceedings, a statement in a document produced by a computer shall not be admissible as evidence of any fact therein unless it is shown:

a) that there are no reasonable grounds for believing that the statement is inaccurate because of improper use of the computer and;

b) that at all times the computer was operating properly or, if not, that any respect in which it was not operating properly or was out of operation was not such as to affect the production of the document or the accuracy of its contents."

This provision requires assurances about the correct use and normal functioning of the computer system at the time the evidence was created. This would require oral or written certification of the reliability of the system used to capture/create the evidence by a person familiar with the system. There is an article in Section 69 that discusses whether computer-generated evidence will be considered direct evidence or hearsay, and the reliability of such evidence. One case discussed in the above article concerns a breathaliser machine – the evidence was ruled as inadmissible because the internal clock in the machine was 13 minutes out. The defence claimed that the machine was not working100% correctly. The Law Commission has recommended the repeal of Section 69 of PACE, which in their opinion "serves no useful purpose".

5.3 Weight of evidence: authenticity and audit trail

Even when a digital image has been submitted as evidence, there is still the problem of proving the accuracy and assessing the weight assigned to such an image. Two areas are important here: the quality of the image, and how well the image can be authenticated. Without authentication of any kind, even an image which clearly shows the defendant committing the crime would have little probative value, and may not even be admitted as evidence. Similarly, an unrecognisable (a poor quality or blurred image) would have little or no weight as evidence no matter how well it could be authenticated.

The authenticity of a digital image may come from several areas:

As with paper records, the necessary degree of authentication may be proved through oral and circumstantial evidence, if available, or via technological features of the system or the record. Oral evidence will concentrate largely on the management and operational procedures applied to the computer record. Ideally these should be assessed by the lawyer in advance to ensure that a satisfactory audit trail can be produced. Circumstantial evidence will cover a wide variety of matters, including the record's consistency with other documents which make up a linked transaction. Technical evidence might come from system logs, particularly if they are designed specifically with this end in mind, or through embedded features of the record itself such as watermarks and digital signatures; although these are only of any real use for authentication if the encryption key is outside the control of the person adducing the evidence.

Digital watermarking is the main way of authenticating digital images, and "can provide a high level of security in conjunction with an audit trail". Digital watermarks come in two varieties, fragile and permanent. Fragile watermarks are those which are destroyed if an image is manipulated in any way - these are useful for proving authenticity. Permanent watermarks on the other hand should be extremely resistant to any form of image manipulation, and are useful for proving ownership and copyright. For the purposes of preventing tampering, fragile watermarks are required, but these are only of value if performed at the time of image capture. Any technology has the potential for being circumvented, but as one witness from IBM said (in the House of Lords report) "*in all of these processes the only thing that you can do is make it very difficult, and if you can make it difficult enough, such as that the process takes too long, then you are at least achieving part of your aim*."

The JISC has published a JTAP Report, Number 34, an "Overview of Watermarks, Fingerprints, and Digital Signatures". While this report is focused on such technology for copyright protection, it provides a very good background into the subject and includes an overview of some current products.

5.4 Digital systems used successfully in court

Several companies now provide digital image capture and storage solutions, and digital signatures appear in many of the specifications. One such company is called Primary Image. In the Q & A section of their "Second Eyes" product Web site, we find the following information:

Q: Will digital images be accepted in court?

A: Yes. Primary Image produced the first digital imagery to be presented and used successfully in a court of law in England. Images have a combination of burnt-in textual annotation and digital watermarking added before being encoded, ensuring tampering is detected.

Loronix, another company that provides digital surveillance equipment, has developed a digital video surveillance system linked to cash register transactions in Dayton Hudson retail stores, and Washington Dulles International Airport. This system captures digital video which is stored on Sony DAT tapes housed in a "jukebox". The recording system is in line with an atomic clock synchronization system that Dayton Hudson uses on its registers, so the time stamp on the video is accurate to within 500 milliseconds. Each recorded event also stores other important information, such as register, number, receipt, etc. Software to search and locate specific events is provided, saving time and effort. Loronix's system tries to ensure that the images captured are tamper-proof by encoding a "fingerprint" into them:

"*Each video clip is fingerprinted through a mathematical algorithm during the video capture process. The fingerprint becomes part of the clip and is used by the playback software to verify the video has not been altered*."

5.5 Conclusion

Digital images may be used as evidence in a court of law regardless of any authentication process, but it is the weight associated with this evidence that will be in question. The two methodologies that will help to establish an image's authenticity are audit trails and technological solutions such as watermarks. Audit trails document any process that the evidence has undergone, its location at all times, and who has had access to the evidence. If the image concerned is captured from a system designed to improve security, then the audit trail should start from the moment of image capture.

Digitally watermarking the image (or using any other form of encryption/tamper detection) as soon as possible after capture (preferably in the camera itself) increases the reliability of this evidence further still. Where images are concerned, we can also stamp other relevant information onto the image as well, such as the exact time it was taken.

Institutions should consider defining policies and procedures for handling video surveillance footage. This may be required by the 1998 DPA anyway, but such a process, describing the audit trail, should also strengthen the use of such data as evidence in a UK court of law.

6 CCTV, Safety, and Lone Working Environments

While the primary reason to install a CCTV system may be to protect equipment from theft and to deter crime on campus buildings, it is also worth recognising that CCTV also has a role to play in both health and safety and lone working.

6.1 Health and Safety

In many situations where students are working in a laboratory environment it is a requirement that students are supervised by a member of staff, such that if an accident occurred the supervisor would be able to take appropriate emergency action. The addition of a CCTV system does not remove the need for direct supervision. However, any extra pair of eyes which may be watching over the laboratory can only improve the chances of an accident being dealt with quickly.

6.2 Lone Working

Staff and students working alone in offices or laboratories are at particular risk if a medical or similar emergency arises. Any system that can monitor lone workers has a value provided that the monitoring is effective. While CCTV in offices may not be cost-effective (or desirable by the staff), cover in laboratories where students may be working on their own (e.g. in a 24-hour workstation room) can again only improve the odds of an accident being acted on quickly.

6.3 CCTV for supervisory purposes

When deploying a CCTV system as an aid to watching over laboratories or similar environments, there is not necessarily the need to use as high a quality camera as when recording video footage with a view to identifying a criminal (or using the footage as evidence in a court case).

It is important that CCTV supervision does not rely on a digitised feed which only activates when motion is detected, as an incapacitated person is not likely to move.

The difference in cost between a high definition colour camera (typically £1,000) and a low definition black and white camera (typically £100) is not insignificant. If budget is an issue, then the higher quality cameras may best be used at entry/exit points in a building, with lower quality cameras deployed in laboratories and other "monitoring" points. Any person seen or recorded on a lower quality camera must have walked past a higher quality entry camera to gain access to the building.

There are other ways to improve security for lone workers. A "panic alarm" can be fitted, or, if a theft alarm is fitted (as described later in this report), certain zones can be "panic button" activated, and the zone information relayed to security staff would indicate the nature of the alarm. Alarms activated by (a loud) voice could also be considered.

Institutions should have a Safety Policy that covers or at least gives guidelines on Lone Working. Certain types of lone working are prohibited by law (e.g. working with live electrical conductors), but other situations will occur on a daily basis, whether inside or outside of "normal working hours". Risk assessment exercises may identify where remote surveillance may be able to contribute to a safer working arrangement, reducing an element of the risk of lone working.

7 Digital Images and Video

This project investigated how to make innovative and cost-effective use of an existing data network (and hosts on that network) to improve campus security. This report discusses the use of video and image transmission over an existing data network as a means of providing a deterrent to potential offenders and also to act as reassurance to those present that they are working in a secure environment.

Several methods have been examined for improving campus security through the use of CCTV cameras, PC-hosted cameras and use of the existing data network and hosts on that network. Such methods have many implications, such as network bandwidth usage, ease of use, security, cost, storage requirements, hardware requirements etc.

Some systems store data to a local PC hosting a camera, while others transmit video data directly over a data network. We have examined methods that utilise dedicated cabling for video transmission, as well as combinations of dedicated cabling and shared data networks.

7.1 Digital CCTV Camera Systems

The most common form of visible security is a set of CCTV cameras placed at strategic locations (most obviously entrances and exits) which then send their video feeds back to a "control room" of some sort. Historically, the standard practice is to record video footage on a time-lapse VCR, and have security officer(s) who can watch the control room monitors, recall the VCR footage, or possibly direct a certain camera's output to a continuous real-time VCR.

This method has two drawbacks; first, there is the need for video tapes. These need to be changed when they complete one tape's worth of recording. Depending on the time-lapse rate, this may be anything from two or three times a day to once a week. The number of video cassettes required for a history trail running back for even a few weeks can be high, and moreover the tapes have a finite life-span, i.e. they can only be re-used a certain number of times before a loss in picture quality results. The second drawback is relying on a human operator to spot when something is happening on one of the CCTV monitors. While door-sensor and passive infra-red (PIR) alarms can alert an operator, it's possible that he or she may be watching one CCTV feed with nothing happening while another not currently being viewed is showing an intrusion.

Digital systems can offer solutions to both these problems. In the first case, the data is stored electronically. The only limit to the amount of data that can be stored is the size of the disk device being used. With PC system hard disks costing little over £100 for 20GB, the potential for mass, cheap digital storage is good. The digital medium does not degrade as the magnetic medium does, and can be re-used indefinitely.

In the second case, digital CCTV systems are able to perform digital image processing techniques such that video clips are only recorded when image differences are detected (typically on a pixel-wise basis). Most will allow the operator to control the threshold of the amount of difference and/or mark an area of the image to look for differences in (for example to isolate a door from a busy corridor).

As a result, when searching for an incident it is probably not necessary to search through hours of recordings that show nothing. With most systems it is possible to search through the incidents

sorted by time and camera, thereby further easing the access to required footage. However, this ability is precisely the one which under the 1984 DPA means your system is capable of automatic processing and as a result is more likely to require registration under the Act (and the 1998 Act).

It should also be noted that by recording only when "incidents" are detected it can be argued that the continuity of evidence is being broken. When managing a tape-based system, it is important to maintain a continuous recording scheme and to avoid breaks in the footage. Doing so, along with logging of the tape usage, helps increase the weight of evidence by reinforcing the audit trail. In the case of digital footage the evidence must be given weight by strong authentication of the time and location from which the footage came. It may be possible to argue that a non-continuous digital video clip omits some important piece of evidence that cannot be seen as it was never recorded. However, as discussed in the previous Chapter on digital evidence, Primary Image has managed to use such a system successfully in a UK court of law.

It is of course possible to record continuous digital footage, provided sufficient local (or remote) disk storage space is available. However, some digital systems that transmit video footage over a data network rely on the "bursty" nature of such transmissions (the fact that surveillance cameras tend to watch over areas with little happening) to allow the video and regular data traffic to co-exist happily without causing network saturation.

A further advantage of digital systems is the ability in most cases to be able to view the CCTV video feeds in real-time from any PC connected to the Local Area Network (LAN) or even in some cases via a dial-up network connection (at a lower quality). The nature of the real-time digitisation may mean the quality degrades, but if the feed is purely to draw attention to an incident, or to allow an operator to turn on high-quality recording elsewhere, that may be good enough.

It is common for CCTV systems to be designed with video switching or mixing in mind. For example, if you have four CCTV feeds and one time-lapse VCR, it may be more cost-effective to use a video switcher or multiplexor rather than buy three more VCR devices.

A typical switcher might allow each of the four inputs to be recorded in a number of ways:

1. All at once, if the switcher can mix four feeds into one (a "quad processor").
2. Each input in turn could be sent to the VCR in a fixed rotation.
3. The inputs could be sent to the VCR in a programmable order, e.g. input 1 for 5 minutes, input 2 for 1 minute, etc.

With a digital system, the video feeds will usually go directly into a multi-way capture card or black box that can replace the video switcher/quad splitter. Most digital CCTV systems can provide continuous recording of each video feed, rather than the switched time-lapse of an analogue VCR system, or can use image difference to record just sequences of potential interest.

7.2 netCam superVisor

The superVisor system from netCam Ltd is advertised as a Digital CCTV Ethernet Network Server. This system was chosen as an example of one that digitises at source and then transmits the footage over a data network. There are two superVisor units available, a single camera version and a four camera version. We tested the four camera version.

Figure 7-1 A photo of the superVisor (from the netCam Web site)

The superVisor system claims to offer the following functionality:

- Video transmission using standard Ethernet local area network (LAN)
- Dozens or even hundreds of cameras can transmit video simultaneously and continuously over the LAN
- Proprietary image compression technique that does not overload the network
- Viewing and recording of video from any PC on the network
- Simultaneous viewing, recording and playback
- Unattended recording capability
- Automatic motion detection

The units are standalone network devices with their own power supply, one LAN connection and either one or four video inputs. Setup was simple, an IP number needs to be allocated for the box (just as per any other Internet device), and the software installed on a PC. The software will only run if a parallel port dongle is attached to the PC.

The superVisor system provides only low resolution video feeds - 384 x 286 pixels, with a contrast of only 128 shades of gray (the video feeds are black and white only, colour cameras may be connected, but the colour will be stripped). This does indeed keep the network bandwidth usage down (64k bits per second per camera connection on average), but means that the images may not be suitable for some applications where it is necessary to clearly identify individuals.

The software is useful and easy to use allowing you to view or record to disk up to nine cameras (per running application). Video streams are recorded to disk in a proprietary format, which can be accessed only through the superVisor software. Video streams are recorded in 15 minute 'files' to allow easy access to a particular time reference, and each frame is time-stamped. Video data is only recorded when motion is detected (and the motion detection sensitivity can be adjusted) which allows a far greater amount of "useful" video to be recorded to disk. For a typical surveillance camera this produces around 500Mb for a 24-hour period. A 20GB IDE PC hard disk drive would thus be able to record four cameras for 10 days. Depending on how long it is necessary to keep the video archive, it is a simple matter to transfer the data files onto a backup device, such as DLT (but note that any such archiving should be documented for purposes of the audit trail).

The motion trigger can also be restricted to a specific area of the frame, for instance a door, whilst ignoring movement elsewhere (people moving around, a road perhaps, etc). Alarms can also be set to detect connection failure, possibly due to a camera being tampered with.

The minimum specification PC for running the superVisor software is a 133 MHz Pentium machine with 32MB of RAM, running Windows. For best performance a Pentium 200Mhz or higher is recommend.

7.3 Primavision Hobo

Another system we tested came from Primavision Ltd, another UK company who deal in digital security solutions. This system requires a PC for the capture station as it is based on a PCI video capture card (alternatively, the system can be bought pre-built, with the software installed). The initial capture card can take up to four BNC video inputs, but is expandable via daughter cards (which do not require PCI or ISA slots), up to 36 video inputs.

Figure 7-2 The Primavision four camera capture card (from the Primavision Web site)

The system's main features are claimed to be:

- Continuous or event based storage from external device inputs or integrated video movement detection with masking
- Expandable from 4 to 32 cameras
- On board fast multiplexing enabling full screen, quad or 9 way viewing, with sequencing
- Simultaneous record and replay, even whilst on-line
- Pan and tilt camera control
- Multi-channel pre-alarm (records images leading up to and following an activation)
- Comprehensive dial-out management controls alarm response

This system is more advanced than the superVisor system, but also more complicated to set up (though still not difficult). The capture card can handle both colour and black and white video streams, and captures images at a resolution of 384x288 pixels. Two different compression levels are available. The incoming video streams are recorded to disk by the camera station (the PC containing the capture card). This can either be a local disk or any available network drive (if the network bandwidth is good enough). Video is only recorded to disk on detection of an "incident". An incident can be triggered by motion detection on the PC, or by connecting the unit to compatible alarm inputs. The motion detection on this system produced less false alarms than on the superVisor system, even when the superVisor sensitivity was set to its lowest setting.

As with the superVisor system, areas of the video may be masked off from the motion detection, but whereas the superVisor only allows a rectangle of interest to be specified, the Hobo system divides the image into small squares, and allows any of the squares to be masked off.

All of the main processing is performed on the PC containing the capture card, and no further hardware is required for normal use. However, if a user on the LAN wishes to view the CCTV video in real time, or review the incident list (and corresponding video footage) from their own PC then the NetView software can be installed to allow this. The NetView software requires TCP/IP networking to be installed, but the computers must be on the same workgroup, and the directory containing the stored video must be available as a shared folder. Note that under the DPA you must take reasonable steps to keep personal data secure, thus the setting up of shared workgroups or data stores must be considered carefully.

7.4 CCTV video switching systems – Microswitcher PLUS

While it is entirely possible to use only solutions such as those provided by Primavision and Netcam, we often need video feeds to go to monitors placed around a building, to time-lapse VCRs (as a back-up, or whilst testing the reliability of the digital solutions), as well as to human operators. This can be accomplished by the use of a video switch, such as that provided by ACI International - the Microswitcher Plus. The basic switch card can be installed into any PC (a Pentium 200 is recommended) and allows for 16 video inputs and 8 video outputs. Expansion cards allow a further 16 inputs per card, up to a maximum of 64 video inputs. Physical alarms can be integrated into the system via a further add-on.

Unlike the Primavision and netCam systems, ACI International is US-based. We experienced no difficulty in ordering a system direct from the USA.

The beauty of the Microswitcher is that it allows any of the 16 video inputs to be mapped on to any of the 8 video outputs (NTSC or PAL). This can be controlled manually at the host PC, manually at a remote PC over the network (via an extra piece of software) or automatically by programming "tours" between inputs and outputs. The software also allows (remote) control of pan, tilt, zoom (PTZ) cameras. The Microswitcher PLUS system also supports a large number of additional programmable features including:

- Define cameras to be displayed on monitors by time and day
- Specify alarm actions by time and day
- Title all cameras with on-screen camera names
- Run video tours by time and day automatically
- Send PTZ-equipped cameras to preset locations automatically by time and day

With such a video switcher it is possible to send four of the eight outputs to a system such as the Primavision Hobo. A fifth output can send a tour of all the inputs to a time-lapse VCR and the remaining outputs could be used to send video to various monitors around the building, which can display tours of the cameras. The Microswitcher can display time and date information on the outputs, which indicates to would-be thieves the fact that a live recording system is in place.

In such a scenario, video is run over dedicated cabling (typically 75 ohm coax, but raw video can also be run over Category 5 "data" cabling – see later in this Chapter). Use of the data network only occurs once the videoswitch directs an output to a PC-based digitising card.

ACI International does sell its own video encoding board, but we have not evaluated it. ACI claims the board can send PAL video at 25 frames a second, using a minimum bandwidth of 640Kbits. We intend to try the add-on alarm device (which can work in tandem with the video switching) in the near future. Our current Microswitcher deployment is discussed in a later Chapter of this report.

The system's only notable failing is that it does not (for US legal reasons, apparently) carry sound, so it cannot be doubled up as a videoconferencing system.

ACI offer free source code to sites wishing to integrate the Microswitcher with other security systems (such as card locks).

Costs: Microswitcher PLUS $2550, or $3049 with client control software (an academic discount is available). Additional slave card (16 inputs) $1250, ACF Alarm board $1191, 16-contact alarm card $157.

7.5 PC Cameras (aka WebCams)

Many PC's come with cameras connected to them, commonly known as WebCams due to their most common use - capturing images that are then displayed on the World Wide Web. The WebCam consists of the hardware (camera) and the software which digitises the camera output. Such systems can be very cheap to acquire and set up for PC owners, with a parallel port or universal serial bus (USB) camera (such as a Quickcam) costing £70 or less. Alternatively you can use a video capture card (such as the Winnov AV) which is installed inside your PC and into which you plug the video output from a (generally higher quality) camera.

We found the PCI-based Winnov AV cards, at around £150 each (US list price $199), provided high quality (640x480 pixel, 16.8M colour) digitisation and would recommend their use. The key criteria for WebCam use are resolution, colour capacity (if required) and speed of operation. Winnov supply Windows drivers (with free updates from their Web site) and there are also Linux drivers freely available on the Internet. One PC can host multiple capture cards.

Two of the better Internet sites for WebCam information are WebCam World and WebCam Resource, which feature useful listings of WebCam hardware and software (they appear biased towards the Webcam32 product, but it is a good product). More recent packages allow WebCam "video", though usually at only a handful of frames per second at relatively low resolution (performance is better when run over just a local LAN).

In the academic community, PC cameras are often used for videoconferencing. However, these cameras lie idle for the majority of the time. One method to provide a little extra security for existing camera owners when absent from their PC is for them to leave it running image differencing WebCam software. Such software works by comparing sequential images and storing them only if there is a significant difference (e.g. when someone is moving in front of the camera). Some systems can store video (e.g. AVI) clips, and some can buffer data so they can save footage prior to the trigger event too. Of course, the last image grabbed may be one of the thief stealing the PC running the WebCam, so it is prudent to store the images on a remote file server, e.g. the Gotcha! product supports image uploading to a remote FTP server (FTP being the standard file transfer mechanism on the Internet).

7.5.1 GOTCHA!

GOTCHA! is a well-specified piece of WebCam software with motion detection and alarm notification designed in; it is specifically aimed at the security-conscious market rather than the more casual WebCam user who wishes to display a single view continuously.

Its main features include:

- Get notified via e-mail with an image attachment when motion is detected.
- Setup GOTCHA! to post (using FTP) images created when motion is detected.
- Get paged when GOTCHA! catches something.
- Call GOTCHA! on a remote PC to view the remote live or recorded video over a phone line.
- Play back the entire day's events (or week's events) in one video time-stamped file using minimal disk storage.
- Alert with audio message when people come into your office while you are away.
- Mask out unwanted movement like traffic outside your window.
- Use a timer to schedule when GOTCHA! should observe. Schedule multiple sessions per day.

Figure 7-3 GOTCHA! in operation. Beware signage…

We tested GOTCHA! in our own project laboratory when we heard that one postgraduate student was having his Polo mints pilfered slowly. GOTCHA! detected the offender the next day. Such use of digital recording is of course subject to the Data Protection Act. In this case the irony is that if the "thief" is known to the individual running the software, he or she can be identified from the image and thus the scenario is more likely to meet the DPA registration requirement.

A multicam version of GOTCHA! allows up to four camera sources to be monitored on one PC, given you have four capture cards in that PC.

Costs: GOTCHA! licence $69.95 (no media), GOTCHA! multicam $399.

7.5.2 Webcam32

One of the most highly recommended WebCam software packages on the Internet is Webcam32. Unlike GOTCHA!, the focus of Webcam32 appears to lie in more mainstream WebCam use, including live chat windows and streaming audio support. It does have some "intrusion" capability though.

Webcam32 supports uploading of images to an FTP server. It also offers Web video "streaming" via either server push (the client browser downloads a JavaCamPush applet to receive images sent from a process which runs on the Webcam32 host) or image pull (using JavaScript to repeatedly fetch images from the Webcam32 host) methods. However, both techniques are inefficient in bandwidth usage as they do not use MPEG-style image difference techniques to only transmit image changes. The frame rate is also likely to be seconds per frame rather than frames per second. The server push can run from a pre-saved AVI file if desired.

Figure 7-4 Webcam32 in action. Lighting is an important issue.

For intrusion detection, Webcam32 has an image compare ability, whereby a percentage change threshold between frames can be specified. If the threshold is exceeded, the image can be either saved locally or uploaded to a specified FTP server.

An official Linux version of Webcam32 is under development at the time of writing.

Costs: Webcam32 licence $25.

7.5.3 I-Spy

I-Spy was the first WebCam software we trialled. Its functionality has been surpassed by Webcam32, and it has no support for intrusion detection applications. One interesting property it does have is support for S/Key one-time passwords for the FTP image uploads.

Costs: I-Spy used to be an independent product but now appears to have been bought by Surveyor, the company who also market Webcam32. At the time of writing both packages can be obtained together for $25US.

7.6 All-in-one WebCams

Some manufacturers produce dedicated WebCams. Such devices are a single compact network device containing a camera and Web server. As a network device it requires an IP address to operate. The advantage of such a system is that it requires no PC to host it, is easy to set up and is rapidly redeployable.

We trialled the Axis Communications network camera, which also has an FTP ability. The performance was good, but the device featured no "smart" ability, i.e. there is no way to use the device to only capture images around one "incident" in the same way as GOTCHA! can. Any "smarts" would have to be implemented at the client end (the process fetching the images from the device's built-in Web server). The bandwidth requirement for the latest model, the AXIS 21000, is around 0.7Mbits for 10 frames per second at 320x240 or 1.3Mbits for 5 frames per second at 640x480.

The major drawback we found with the camera at the time of testing was the lack of a wide-angle lens. In most surveillance positions, a wide-angle view is desirable. However, the AXIS 2100, costing around £800, now features a replaceable CS-mount lens. Another concern is that the built-in Sony camera doesn't work down to as low a light levels as it might ideally do. If poor light is a concern, a dedicated camera capable of lower lux levels may be preferable.

7.7 Remotely viewing CCTV video

There are a number of software products designed for video transmission or streaming.

Our initial pilot scheme in the Department used CU-SeeMe to allow any client to view footage from the CCTV cameras. CU-SeeMe is a videoconferencing tool that allows one to one, one to many, or many to many videoconferencing. One to one is simple as two clients directly connect to each other, but one to many and many to many require the use of a reflector. The reflector is basically the meeting point. Every user connects to the reflector, and can see all others connected. The reflector gives control over who can broadcast video. To use CU-SeeMe for viewing CCTV video in real-time, we simply connect a CCTV feed into a PC with capture card running CU-SeeMe. The capture PC sends its output to the reflector, from where any user running the CU-SeeMe client can view the video. Bandwidth used for a discernable image is of the order of 100-200Kbits.

Unfortunately the CU-SeeMe reflector software is no longer freely available (the new product, MeetingPoint, is rather expensive) though an unsupported black-and-white only Linux reflector can be found on the Internet. As a result, CU-SeeMe is no longer a particularly cost-effective way to distribute CCTV footage.

We trialled the Real Video software from Real Networks. Our Department had the server software licence already available. The quality of the Real Video feed can be adjusted in the client we ran by raising the maximum bandwidth used up to a limit of 500Kbits, which gave reasonable quality video but certainly not of broadcast quality. A bigger concern was the transmission delay. Using a PC to generate a live Real Video stream we saw delays of up to 10-15 seconds between the capture and viewing of the video stream. In many applications this would not be important, but if the software is being used (for example) to control a PTZ camera,

the feedback delay would cause significant problems. And any intruder would have something of a "head start" on any security staff.

7.8 Bandwidth usage and Quality of Service

The bandwidth consumed by running video applications over a data network must be considered carefully before deploying a system, to avoid a scenario whereby "regular" network users do not have their service adversely affected by the video applications.

The two sections below briefly discuss Quality of Service (QoS) methods for Ethernet and routed IP networks. In a campus or building security scenario, most of the data transmission would be over Ethernet, with a limited number of routers on campus perhaps being used. It is unlikely that CCTV footage would need to be transmitted between JANET Institutions, but that doesn't necessarily mean that Institutions shouldn't require QoS methods to be supported across the JANET network.

There is a very good JISC/JTAP Report, Number 36, titled "Network Delivery of High Quality MPEG-2 Digital Video" in which methods for delivery of quality video streams are discussed. MPEG-2 is the technology used in DVD, and though the bandwidth requirements of various "flavours" of MPEG-2 vary, it will typically require 4-6Mbits of bandwidth.

It is of course possible to run video over data networks using TCP/IP video applications, but to do so over a dedicated private data network. This also has the advantage that the data is less likely to be snooped by regular network users (a property which is desirable under the DPA).

7.8.1 Quality of Service over Ethernet

One of the problems with attempting to transmit broadcast-quality video over an Ethernet network is the general lack of Quality of Service support. The current standards favoured for QoS over Ethernet are IEEE 802.1Q and 802.1p. The former defines a method for offering virtual local area networks (VLANs), the latter a MAC (link) layer method for prioritising packets to one of 8 priorities.

The unfortunate property of Ethernet networks is that saturation occurs at around 40% of theoretical bandwidth capacity in a multi-host environment due to the statistical properties of the CSMA/CD algorithm such networks use. A point-to-point link in a 2-host network will approach the theoretical limit, but such scenarios are atypical. Thus a 10Mbit Ethernet is unlikely to be able to deliver a 4-6Mbit MPEG-2 stream on top of regular data services, while a 100Mbit Ethernet network may be able to cope with up to 6-8 high quality streams, if the network is dedicated to streaming. The advent of Gigabit Ethernet does offer the potential for combined data and streaming applications to co-exist more readily, though at present the costs of Gigabit Ethernet remain relatively high. Also, many Institutions may only be considering upgrading to 100Mbit networks, yet alone Gigabit.

7.8.2 Quality of Service over Routed Networks

At the Internet Protocol (IP) layer, there are two main methods for offering QoS. The first, Integrated Services, or IntServ, works by attempting to reserve a certain QoS (generally a minimum bandwidth path) between all routers that lie on the network path between two hosts

(most likely at different sites if separated by more than one router). The Resource Reservation Protocol (RSVP) is used to perform the QoS negotiation. IntServ is well-suited to video transmissions, but router support is variable, which given the fact that the QoS obtained is only as good as the "weakest link" is often somewhat problematic.

The other router-resident QoS method is Differentiated Services, or DiffServ, which works more like 802.1p by assigning classes (priorities), in this case to IP packets. There are four main classes, the highest being expedited. The problem with DiffServ is that, for political reasons, it is usually only applied across the network of one Internet Service Provider (ISP). If the JANET Academic Network adopted DiffServ, it could apply packet prioritisation across its entire educational network, given router support and co-operation from the regional MANs.

7.9 Video over dedicated cable (including twisted pair)

Due to the difficulty in providing good QoS on Ethernet networks, it is currently still very difficult to offer broadcast-quality video to the desktop. While a handful of medium quality streams may be deliverable over a lightly-loaded data network, the quality will not be equal to that of raw video over regular 75ohm coaxial cable. The video may be sufficient for videoconferencing, but be less than ideal for CCTV recording and use in a court of law (though it may be enough to alert a viewing staff or security person to a potential incident).

The best quality imagery is still to be obtained by running CCTV or video over dedicated coaxial or fibre-optic cabling. However, it is also possible to send near broadcast-quality video and stereo audio over Category 5 twisted pair cabling using a system called Sourcerer.

The Sourcerer system Web site explains their technology:

The Sourcerer Structured Cable Television System uses state of the art technology to send and receive broadcast-quality video and audio along twisted pair data cables. It is an active system resulting in signal quality which is virtually unchanged even after transmission over kilometres of cable. It uses existing structured cabling to distribute terrestrial, satellite and cable televison channels and indeed any composite video source such as TV cameras and VCRs. It works with domestic televisions, monitors and personal computers and each user can independently select and view channels using an infra red remote control keypad

It is rare in modern offices to have RF (radio frequency) coaxial cable and aerial sockets installed as a matter of course. In most new office buildings it is standard practice to install low-cost multi-pair data cable - often referred to as structured cabling, or UTP, FTP or CAT-5 - for use in computer networks and telephone systems. The Sourcerer system can be patched-in to provide TV distribution using spare cables already installed in the building.

The Sourcerer system can be used in video distribution applications such as airport information systems, point-of-sale promotional video and presentation message systems, where a number of monitors display the same signal source.

Such a solution would require a site either over-specifying the number of required network points in a new building, or running extra Category 5 (data) outlets, so that there exist spare points to connect the Sourcerer system to. These points could be used not only for CCTV cameras, but also for videoconferencing, remote viewing of lectures, or any other video application (because stereo sound is integrated into the product). The receivers are small and easily portable, allowing for a very flexible solution whereby ad-hoc surveillance can potentially be performed at any data

outlet, given the building is flood-wired in such a way that that outlet can be patched back to the receiving Sourcerer box.

The cost is around £200 each for the transmitter and receiver. Thus a redeployable surveillance capability running over one data point would be £400. A two-way videoconferencing ability running over two data points would cost around £800 (plus the cost of two cameras, microphones, speakers and monitors).

7.10 Conclusion

While it would be ideal to be able to make heavy use of local Ethernet networks for high quality video transmission, the technology to enable such an application is not quite available yet. The advent of Gigabit networking will facilitate such applications, in particular when combined with the IEEE 802.1p standard for packet prioritisation. Both technologies are still in their relative infancy (e.g. Gigabit was only standardised in late 1998). In a routed IP environment, QoS support needs to be offered in the routers, via IntServ (RSVP) or DiffServ, and this QoS has to be seamlessly integrated with the Ethernet QoS method.

Until QoS support in data networks is widely available, there are two options to consider:

1. Using dedicated cabling for video applications (be it data or raw video). Devices such as the Microswitcher PLUS can be used to automatically or manually switch selected video inputs to outputs which can perform intrusion detection (e.g. via GOTCHA!) and which store the resultant digital data. By concentrating video sources at a video switch, outputs can be sent to a variety of devices, be they traditional VCRs, quad processors or digital video or image capture devices.

2. Running off-the-shelf lower quality video applications over data cabling. The netCam superVisor delivers lower quality video (than "native" CCTV), but will operate reasonably well over a 100Mbit network. If only a few camera positions need to be monitored, and image quality is not a prime concern, such a solution may perform adequately well.

In the subsequent Chapter on our own deployment experience in the Department, we illustrate and discuss our experiences having chosen to follow the first option (though our path deviated somewhat as we progressed).

8 Social aspects of CCTV deployment

The full social impact of the deployment of a CCTV system is beyond the scope of this project. However, in the course of our project we did receive a number of comments, though all were made "informally" rather than being "official" complaints.

In response to suggestions that CCTV would help improve staff and student safety, one member of staff replied:

"*I don't think the majority of people "feel safer", the type of intruders we get here are not bothered by security cameras, however some of the legitimate inhabitants feel spied on… I think people will need reassuring that the cameras are not going to be used "in evidence against them". I know someone who works in a travel agency where they have this so-called security camera - it turns out it is also being used to watch the staff while they work and censure them for time wasting."*

This comment was typical of four or five comments received (out of 200+ staff in the Department). There has been similar concern expressed over the use of card lock access logs, which would be a greater indicator of staff movements were card usage required in regular work hours instead of just outside work hours. Likewise, another security measure, the current ID card system (whereby all staff are obliged to wear their ID cards when in the building) was not universally well-received.

The Data Protection Act is, in theory, aimed at protecting the rights of individuals when their movements are recorded "digitally". How effectively the DPA is policed will be an interesting issue, but Institutions deploying CCTV and/or card lock systems cannot afford to do so without investigating the implications of the new DPA first. The new Act significantly improves the rights of individuals, both in terms of rights of access to their personal data and in what Institutions can do with that data.

CCTV cameras are becoming more and more common, from car parks, shops, tube stations, city centres, and now in the workplace. The explosion in the use and visibility of CCTV systems could have considerable implications for our civil liberties. There are many concerns over the use of CCTV systems, such as invasion of privacy, reduction of frontline policing and that of crime simply being displaced.

One of the properties of the Internet is that it makes everyone a publisher, and gives everyone a platform to "speak". It is thus no surprise that there are many Internet sites, both "amateur" and professional, which focus on privacy-type issues. These include Privacy International, Statewatch, Lycos Surveillance and Watching Them Watching Us (see the References at the end of this report).

To help ease concerns, any CCTV deployment, whether conventional or digital, should be preceded by explanationary literature for all staff and students affected by it.

9 Deployment of a versatile video surveillance system

This Chapter describes the experiences we went through in deploying a CCTV system in our own Department. The process began shortly before the JISC funding was awarded, and has been ongoing since it ended. We see the issue of building security (much like network security) as being one subject to constant review and refinement.

9.1 Requirements

The main impetus to deployment came through concerns over the level of thefts from Department property. Two instances in particular had severe repercussions as in each case a number of PCs were stolen and, despite a central backup system being in operation, some data valuable to individuals was lost. Walk-in thefts are also a problem, with laptop PC's being the prime target.

Given that a number of cameras are to be installed at a site, there are a number of obvious uses of the equipment:

- Deterrence of theft or crime
- An aid to prosecution of an individual when a crime has occurred
- Improvements to safety and lone working environments
- Videoconferencing

Our main reason for installing a CCTV system was deterrence of theft, and that was the basis upon which we progressed, though we continued to consider the other uses.

Given that we had control over the wiring in New Zepler, we decided that we should seek to utilise a variety of media for video transmission:

- Dedicated 75ohm coaxial video cabling
- Video through software over TCP/IP (Internet Protocols) running on Ethernet or ATM
- Video over Category 5+ unshielded twisted pair (UTP) data cables

The current standard wiring used for new data network installations is Category 5+ cabling, which is able to support 100Mbit Ethernet networking as well as a variety of other services (e.g. voice, ISDN and PC serial links).

New Zepler is a four story building. We specified that each floor should be flood-wired with Category 5+ cabling, with all cables on each floor running back to a patch rack in a single data control room on each floor. These rooms were placed to align vertically such that a data riser space could be used to route cables between the floors. Our data network backbone was ATM over fibre, but we specified extra Category 5+ patch cables to run between floors. We also, as described in the next Chapter, specified extra Category 5+ cabling for the physical security alarm system. We felt that over-specifying Category 5+ cabling was not a problem, as that gave us the flexibility to relocate equipment or to utilise the cabling for non-data purposes. The major cost of installation is the labour element, so choosing to run more cabling at a later date (as we felt would be inevitable) would prove to be relatively expensive.

We also chose to install 75ohm coaxial video cable on each floor. The cables were run back to the data control room on each floor, and from there patched down to the ground floor data control room where we chose to place our main video rack. The data, video, alarm and power cabling were each routed down a separate channel in the cable trunking.

9.2 Considerations for camera placement and technology

Given our main aim of deterrence of theft, our priority was to locate cameras at all entry/exit doors in the building (both New Zepler and the existing linked Mountbatten building which our Department also occupies).

Deterrence requires visual feedback that a system is live, and thus we also chose to locate a monitor at each entry/exit door covered by a camera. The monitor should display either a tour of the views from the other doors or a "quad" picture showing multiple doors. Ideally the entry/exit name and current date and time should also be displayed. It should be noted that under the 1998 DPA there is now a question mark over whether such visual feedback is appropriate. Registration will be required, and as a result safeguarding digital or VCR recordings will be an issue of importance.

Ideally we would deploy high quality, high definition cameras where needed throughout the building. However, we recognised that to do so would be a costly exercise. As a result, if a camera was to be deployed for a monitoring purpose (an aid to safety or lone working) we would consider deploying a lower quality black-and-white camera. Where identification was required, such as at doors, we chose to deploy high quality cameras from RS, colour where possible. A high resolution, black-and-white camera capable of delivering an image at 640x480 resolution costs upwards of £300, plus £100 or more for the lens. All monitors were compact 9" black-and-white Maplin models.

Cameras and monitors were wall- or ceiling-mounted, and powered from always-on fused power outlets. The lux rating for the camera is important; the lower the rated lux level the more able it is to offer a discernable image in low light conditions. Lenses were chosen to give an appropriate field of view for the placement; in a corridor a 30 or even 15 degree field of view is well-suited, while in the corner of a room or overlooking a door a 60 or perhaps 90 degree field of view is required. Cameras with built-in lenses can be bought, often quite cheaply, but the image quality is unlikely to match the dual component systems. Where possible the camera and cables were mounted in a tamper-resistant metal box, and physical alarms should be considered for the cameras themselves.

Even though the location of a camera may have been determined, the positioning is also important. A decision needs to be made on whether you wish to catch an "intruder" on the way in or out of a building, or both. There should obviously be no obstructions to the camera view. Lighting should be considered, both in terms of offering enough light to make any images useful (in combination with an adequate lux rating) and in terms of interference in image quality through sunlight. The direction of the sunlight will change through the day; potential camera positions should be studied for sunlight patterns for all hours of daylight, all seasons of the year. A silhouetted face is unlikely to be recognisable, even with post-processing of the image.

9.3 Designing a versatile system

As described in a previous Chapter, there is a wealth of digital video/CCTV solutions available. However, rather than run a purely digital system we wished to design a solution which was flexible and versatile enough to operate with and over multiple media. We felt that we should also experiment with a variety of techniques such that we could report on their effectiveness in this project.

The main components we now use in our live operational system, beyond the cameras and monitors, are as follows:

9.3.1 Video rack

The heart of our coaxial video cabling, we have a full-height 19" rack in our ground floor data control room in which the majority of our video systems is housed. At the foot of the rack are four BNC patch strips into which the coaxial cables from each floor are fed. We are then free to patch video between any coaxial outlets in the building (of which there are some 50 or so) or into any of our video system components. The outlets can be used as either inputs (cameras) or outputs (monitors or digital capture devices).

9.3.2 Quad processor

To be able to display four video feeds on one monitor a quad processor device is required. The device we selected is able to sit in series with up to four coaxial video feeds and generate two quad video outputs from those feeds. It is thus possible to generate a quad image to be recorded to a VCR while feeding each camera output on from the quad processor to another device (e.g. a monitor or a PC capture card).

We chose a quad processor that could enlarge each quadrant of a recorded quad image on playback, and which, as an option, was able to add caption and time/date information.When digitising quad processor output, the resolution of the system is an important issue to consider, as you will be running at a quarter of the image size of a full screen (and thus the ability to recognise individuals will be reduced).

9.3.3 Time-lapse VCR

While we knew we would be recording digital footage from our cameras, we also wanted to be able to use traditional recording techniques as an extra layer of security. We thus bought a time-lapse VCR capable of recording output from either one camera or our quad processor over an extended period of time.

Such VCRs vary in record time from 4-8 hours (standard live VCR), through pseudo "real-time" devices which record 8 frames per second (and thus which run for 24 hours on one tape), to much longer-term VCRs which record for up to a week (at a frame per second or less).

9.3.4 PC with Microswitcher PLUS card

Given that we had the ability to manually patch between any of our coaxial video points in our video rack, it was a natural extension to be able to want to have control of the patching (switching) via software on a networked PC.

The most cost-effective switching system on the market that met our needs was the ACI International Microswitcher PLUS. This card fits inside a PC, takes 16 BNC video inputs and offers 8 BNC video outputs.

Figure 9-1 The videoswitch PC, also running I-Spy

With 16 video inputs it is possible to feed camera inputs through a quad processor into the Microswitcher, and to also feed one output from the quad processor back into the switcher. Programming tours is very easy to do; any input can be mapped to any output for any period of time in any sequence. Programming can be done at the host PC or remotely. The remote software could be either ACI's client software, or it is also possible to use VNC to operate the hosting PC over the network (though we currently use the client software). Either way, a four-digit PIN number is required to gain control of the switcher.

The remote control of the switcher (which includes the ability for PTZ control) can be offered either locally or to dedicated security staff in another part of a campus. With one switcher output directed at a video streaming package, the remote user can choose which of the (16 or more) inputs to view on their local PC screen.

It is also possible to deploy multiple Microswitchers, and have the inputs/outputs mapped across the multiple devices. When driving 75ohm coaxial cable video, ACI suggest that the maximum cable run should be no more than 300 metres. We ran some Microswitcher feeds to over 150 metres without any observed problem or picture quality loss.

We were also able to host video capture cards in the same PC as the Microswitcher card. One output of the switcher could thus be fed into (for example) a resident Winnov AV capture card. With a copy of GOTCHA! running on the PC we were able to record digital video clips of incidents locally, or upload snapshot "intrusion" images to a remote FTP server.

9.3.5 Sourcerer

We used Sourcerer transmitter and receiver devices to enable us to patch a camera into our videoswitch system from any data outlet in the building. While the cost of the transmitter and receiver was not small, the flexibility offered made the combination worthy of purchase. The cost is also more competitive when viewed against devices such as the superVisor.

9.3.6 GOTCHA!

Of all the low-cost "intrusion" WebCam software we tested, GOTCHA! appeared to perform the most satisfactorily. Its functionality is very good; unlike many WebCam packages it is designed for surveillance use. However, while it will time-stamp images and video clips, it does not use any "fingerprinting" technique. To increase the weight of evidence of such data, additional authentication techniques need to be applied to the acquired video images and clips.

It is also unfortunate that GOTCHA! uses a proprietary file format. A program (gview) is included for viewing the captured video on any PC. When running, GOTCHA! can dump each hour of observed activity to a single file, with the file name generated from the date and time. This file can then be copied or uploaded to another server for future reference (ideally after being "fingerprinted" in some way). An hour of data from surveillance of a relatively quiet foyer area at night generates around 3-5MB of data per hour, at 640x480 resolution. During regular work hours that volume can rise ten-fold or more. Given GOTCHA! stores frames in memory (for speed), RAM is as important an issue as hard disk.

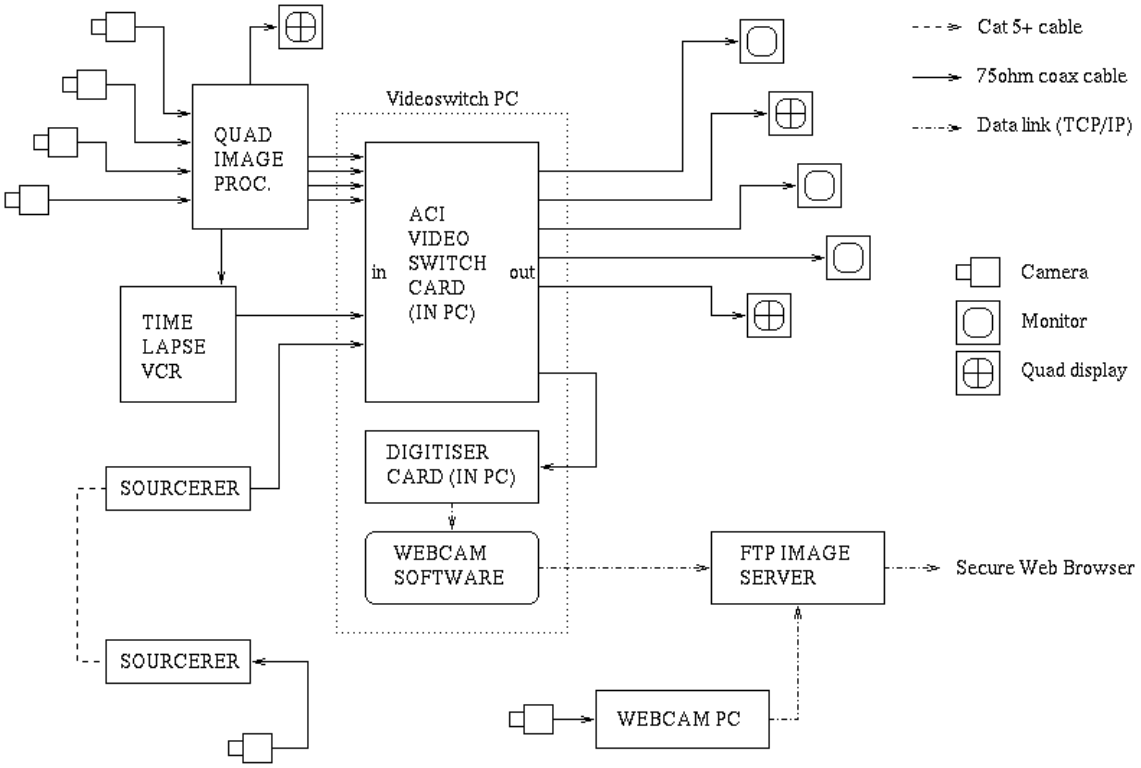

Figure 9-2 A versatile surveillance deployment

9.3.7 Webcam32

General-purpose WebCam software such as Webcam32 may not be quite as sophisticated as products like GOTCHA! for intrusion detection, but will be able to upload images (perhaps in JPG format) to an FTP server if left running while the PC/workstation is unattended.

By offering users file space on a central secure FTP server, they are able to leave their WebCam running in an "intrusion" mode while away from their PC. Such a voluntary system can potentially be of use should an incident occur, though the weight of evidence may be less than a more formal procedure. By using a remote FTP server, an image will not be lost should the incident be a thief stealing the WebCam PC itself.

Many packages which upload images via FTP will use the same filename on each upload; in such cases it is necessary to have a daemon process of some sort running on the FTP server which renames or copies the images to a safe location.

9.3.8 NTP server

It is very important when attaching time-stamps to files, or even just displaying the time on video outputs that are being VCR-recorded, that the time used is accurate and correct.

There is a standard Internet protocol for synchronising time over a network, namely the Network Time Protocol (NTP). By running a client program on a workstation or PC, an accurate source of time can be contacted to allow the local clock to be corrected, if necessary. NTP allows for network delays in exchanging packets, and has safeguards against failures of the remote NTP server.

Figure 9-3 Real-time VCR, Sourcerer, Quad processor, video PC and coax patch panel

JANET run a NTP service which could be used by security applications seeking accurate time. There are also many other NTP servers around the Internet. If you fear local network outages then you can run NTP from a local device with accurate time; many network routers (e.g. from

Cisco) can serve NTP. Cheap "Rugby clocks" can also be obtained, or for a high-end solution (£2,000 or over) you can buy a system that will co-ordinate time via GPS satellites.

We looked at a handful of NTP clients. Of the freeware Windows products, we found that the Dimension 4 client offered everything we needed. Such a client doesn't require much functionality, but editing the server list and setting the synchronisation periods was very straightforward. We have adopted the client on all our security-related PCs.

9.4 Costs

The following table of costs illustrates pricing at the time the project was undertaken. Prices will obviously vary with time. When buying a product you may wish to consider balancing the price against the desirable features we have mentioned in this report.

The cost of installing coaxial video cable will depend on the nature of the building, the quality of cable and the labour element required. It is certainly likely to exceed the cost of installing Category 5+ data outlets (which would currently cost in excess of £50 per point, again depending on the specific situation).

Microswitcher PLUS, with client software	$US 2550
Video rack (19"), 4 BNC patch strips, 3 rack shelves	£ 1,400
50 (approx) coaxial video outlets run to video rack	(variable)
Time-lapse VCR	£ 700
Quad processor	£ 500
PC to host Microswitcher	£ 600
FTP video image server (running Linux)	£ 1,000
Sourcerer transmitter and receiver	£ 400
High resolution black-and-white camera, with lens (each)	£ 420
9" black-and-white monitor, with mounting bracket (each)	£ 115
Winnov AV PCI capture card (each)	£ 120
Licence to run GOTCHA! (each)	$US 70

9.5 Policies and procedures

The prime considerations for policy lie in the requirements of the DPA and in the desire to gather footage that would have a high weight of evidence if used in a UK court of law. Naturally, these considerations should be made before deployment, rather than retrospectively.

For the purposes of the DPA, the main requirement is to first decide whether your system requires registration under the Act, and, assuming it does, to consult your Institution's Data Protection Officer on how best to proceed. It is likely that an existing registration can be extended. Once registered, you are then obliged to certain actions. You will need to be able to produce copies of personal data on request. You will need to specify who handles the data. You will need to take reasonable precautions to secure the data, and certainly to consider the new eighth principle concerning transmission of the data outside EU countries.

For production of evidence, there is some overlap with the DPA. You will need to document an audit trail for data. This should at least describe where data is gathered from, what software and hardware it is gathered by, where the data is stored, and who handles or is responsible for it. Ideally some "fingerprinting" system will "stamp" the data for authentication purposes (we are at

present working on such a system, with generic application, in the Department). There is no legal requirement (as with the DPA) to follow procedures and policies to increase the weight of evidence of any digital data gathered, but doing so may prove valuable in the longer term. Similar authentication techniques may, for example, be equally applicable to video data as other data (such as Unix system logs).

9.6 Data security considerations

It would be possible to write a separate report (or book) on the subject of computer system security. When handling video data there are many general precautions that can be observed. These include:

1. *Securing the system on which data is collected* (the PC hosting the capture or switcher card). There are general procedures for Windows and Unix machines which are beyond the scope of this document (see, for example, Practical Unix Security by O'Reilly). The use of a firewall may be appropriate (see the forthcoming JTAP-631 Report). Physcial access to the system should not be overlooked.

2. *Securing data in transit*. This may be "tapping" done on coaxial video cables, or, on a shared data network, "packet sniffers" may gather data being sent between hosts. One technique to reduce the risk is to use a separate private data network for the video data; this also reduces the contention with normal data traffic. The capture card PC's and any "monitoring" PC can run on a network not connected to the regular network (or Internet). Data served from a web server can be encrypted via use of an SSL certificate (e.g. from Thawte or Verisign) on that server. Data being uploaded can be made more secure via use of (for example) scp (secure copy) rather than FTP.

3. *Securing data while stored*. Products are available (e.g. the F-secure suite from DataFellows) that can encrypt stored files. One should also ensure that if data is being stored on a shared network drive (under Windows, for example) that the drive is protected by at least some password authentication.

The Department is currently investigating secure Internet protocols under a JISC-funded project, JTAP-659, "Secure Internet Protocols for the HE Community". An interim report was published in the summer of 1999.

9.7 Conclusion

Ideally, all video surveillance operations in a typical CCTV deployment could be run over an existing data network. However, we found that products on offer which did digitise video at source (such a superVisor) did so in a way that reduced the image quality to below what might be expected of traditional VCR techniques. Higher-end equipment (typically ATM-based) which allows for broadcast-quality video over ATM will currently cost a significant five-figure sum and is this not suited for a smaller scale deployment.

At a campus-wide level, running up to 50-70 4-6Mbit MPEG-2 quality video feeds over a data network would require a bandwidth at the core of the network of 200-400Mbps, a figure beyond OC-3 ATM (155Mbps) and one which stretches OC-12 ATM (622Mbits). For this reason, the campus-wide CCTV system deployed at the University of Southampton last year was run over dedicated video fibres, albeit at a cost of at least £150,000.

We feel that the most cost-effective building-level option at present is to deploy a versatile system that is a blend of traditional video feeds and data network-enabled applications. By using a video switch device which can be controlled over a data network, and where the video outputs can be recorded traditionally to VCR or digitised to single frames or "incident" clips, we have an effective hybrid system which makes best use of existing and new technologies.

By offering an image upload FTP server to users of the local data network, any PC or workstation with a WebCam capability can be used to gather images when the user is away from their PC. In such situations a cheap WebCam package such as Webcam32 is ideal. For more dedicated surveillance purposes, software such as GOTCHA! is well-suited. For protection of a small building or perhaps just one entry point, a PC running GOTCHA! is a perfectly good option.

The deployment process should not overlook the non-technical issues. These include considerations (procedures and policies) to meet the requirements of the 1998 Data Protection Act and to improve the weight of evidence of any data that is gathered.

10 Use of twisted-pair data cabling for anti-theft alarms

The presence of a visible anti-theft alarm system is an important aspect in deterring either casual walk-in thefts or more determined "targeted" raids.

The two main types of alarm system are:

1. *Motion detection systems*. Typified by infra-red "burglar" alarms. While practical in an area not accessed overnight, deployment is, for example, a computer laboratory that is open 24 hours a day is not so well-suited. Likewise for a regularly-visited entrance foyer.

2. *Physical alarm systems*. Items are secured by wire, breaking the wire sets off the alarm. While such systems are not well-suited to protect laptop PCs, they provide an excellent general deterrent to an opportunist thief.

There are many such systems on the market, all effective in what they do, but invariably requiring specialised cabling and devices to be installed.

10.1 Physical alarms running over Cat3/5 cable

The traditional method for running security alarms involves using dedicated alarm circuitry which can be both expensive and also very inflexible, in that the cabling is dedicated to being alarm-only. The system deployed as part of this project is a blend of two relatively novel techniques. The first is the use of intelligent Device (iD) chips to simplify the alarm circuitry, the second is the use of standard Category 3 or Category 5 cabling to run the alarm signals over.

The iD chip (or "biscuit") is a component developed by a company called Advanced Design Electronics. It allows each chip to have one of up to 30 different "addresses" such that when multiple sensors are wired back to a single alarm panel the panel can immediately identify which of the sensors the alarm originates from. The iD chip resides in a logic board that is mounted in close proximity to where the sensor(s) will be attached. Each iD chip becomes one zone for alarm identification purposes, but multiple sensors can be attached to one logic board and thus be in the same zone.

The wiring used is Category 5 cabling, but the system has also been tested over Category 3 cabling. This is attractive because it means that an alarm system can be installed over existing Category 3 data cabling if the site is upgrading its data cabling to Category 5 or Category 5+ (the higher Category cable is more readily able to run 100Mbit Ethernet networking, whereas the Category 3 cabling is only rated for 10Mbit Ethernet). Thus a site looking to improve network performance and install a reliable alarm system could consider re-using its existing data cabling, thus saving on costs.

In installing such a system in the Department, we worked with a compnay called IT Guard, for whom we are now developing additional security-related software (as explained later in this Chapter).

10.2 Alarm system description

The design of the system we installed is relatively innovative in the way it combines the iD busicuit devices with the use of twisted pair "data" cabling.

The system is centred around a number of alarm panels (up to four in the current design). There is one master panel, and each panel can handle up to 30 iD chips, or zones. There are thus up to 120 zones available to deploy. Each panel in effect adds a unique prefix to the iD chip identifier, so every zone in the building is able to have a unique identifying address. As New Zepler is a four-story building, the obvious choice was to place one panel on each floor. The arrangement of laboratory and office space allowed us to have one zone per office and (approximately) one zone per six desks in the laboratories.

Each zone has a logic board, from which we run alarm cables to the objects we wish to secure. The cables are Category 5+ twisted pair, and to distinguish them from the regular data network we used red cables for the alarms. The cables connect to sensors on the alarmed items, and have two twisted pair sockets, allowing them to be used in series, with the last item having just one alarm cable in and a terminating plug on the other socket. Any break in the cabling will trigger the alarm, though the master panel will only identify the zone, not the exact item being removed.

Each logic board is patched back (just as per regular data cabling) to a 19" rack in which there is a patch strip from which alarm points are connected to the alarm panel. The master panel is able to drive siren alarms, and also communicates all events via a serial link (which also can be run over Category 5 cabling) to the system PC.

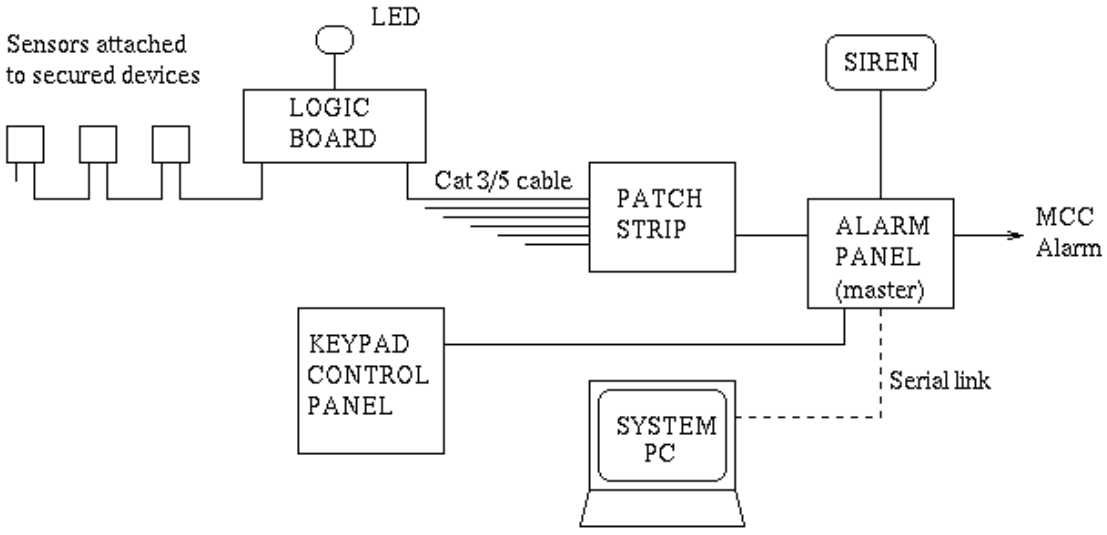

Figure 10-1 Illustration of the alarm system

We installed one keypad control panel per alarm panel. The control panels (keyboard interfaces) can be attached to alarm panels allowing staff the ability to perform such tasks as acknowledging or canceling alarms, or deactivating a zone while a piece of equipment is being moved.

Figure 10-2 The IT Guard alarm panel patch cables

10.3 Component development

New Zepler was wired for physical security by a company called IT Guard, though the complete system came from a number of sources:

1. *Alarm sensors*. The sensors attached to the items to be secured were developed within the Department. This was done as the in-house sensors were both cheaper (working out to around £10 per sensor rather than £30) and also less bulky to install. The alarm sensors we developed in-house are primarily designed to be used on PC casings, but can be applied to any screw fitting. The sensor consists of two circular washer contacts mounted on a small PCB that are held together by the pressure of the screw being in place.

2. *Alarm wiring*. The wiring was installed as part of the New Zepler project. The wiring specification was simply made in addition to the number of Category 5 outlets desired for data (and other) purposes. The alarm outlet density was 1 per office and 1 per six workstations in a laboratory area.

3. *Alarm logic boards and panels*. These were supplied and installed by IT Guard.

4. *Alarm software*. IT Guard installed a monitoring PC. However, because the Department wished to utilise the alarm event logging in customised applications, the monitoring system was extended by tapping into the serial link from the master alarm panel to the monitoring PC. This enabled us to, for example, display alarm events on the building foyer's rolling Web display.

10.4 Software development

The IT Guard monitoring PC collates and logs a wide variety of alarm events. In addition to theft alarms, such events as zone activation/deactivation and under- or over-voltage readings from alarm panel power supply units are reported.

To be able to record and act on the events, we tapped the serial connection to the monitoring PC and fed the signal into a second PC (the "event monitor PC"), where the events were recorded to a database. We could have used the IT Guard PC and tapped the signal to the second serial port, but it was felt that any experimental or development work should be done on an isolated PC.

Figure 10-3 The in-house alarm sensor

The master alarm panel has an output which is fed to the University's Maintenance Control Centre (MCC). When the alarm is triggered, there is a short delay (a period of time in minutes which is user-definable) before any audible alarm is activated on site. An alarm notification is still sent to both MCC and to the defined local control panel. The delay allows "false" alarms caused by, for example, students in laboratories to not cause disruption to other people at work. It also gives local staff who monitor the control panels a good chance of apprehending a real thief who may be unaware that an alarm has been triggered.

If MCC dispatch a member of security staff in response to an alarm, it is very useful for the security guard to know exactly where the alarm has been triggered. For that reason we chose to develop software, in partnership with IT Guard. Which would allow more detailed alarm notifications to be sent across a TCP/IP network (locally or to, for example, a PC in the MCC). For local notification, we chose one display to be our foyer rolling Web display, so that any staff present when an audible alarm has been triggered can see exactly which zone is affected.

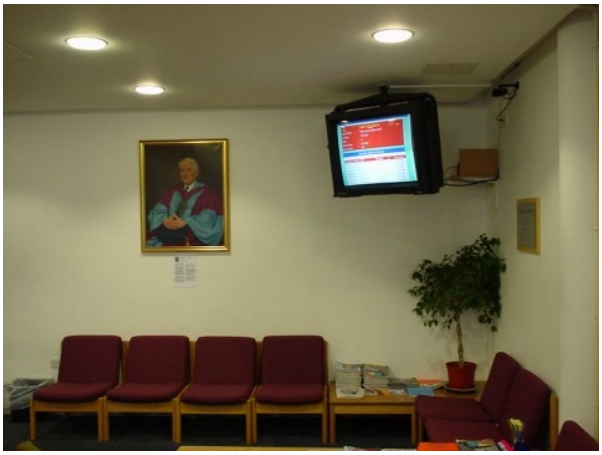

Figure 10-4 The rolling Web display in alarm event display mode

The event monitor PC runs a Visual Basic program that injects the alarm events into a database. It uses the Microsoft Network Dynamic Data Exchange (NetDDE) protocol to notify the rolling Web page PC of a theft alarm. Another Visual Basic program running on the rolling Web PC

listens for the NetDDE call, and then switches the rolling display to a static alarm display. The static alarm page lists the current alarm zone, and the previous eight alarm events.

Through having built a database of zone to room relationships, we are able to list the room number corresponding to the zone in the alarm display. The rolling Web page is one example of the use of a network connection to integrate the "traditional" alarm system into a data network. We are also offering the same data to MCC, so that they know exactly which room they are dispatching a security officer to when there is an incident. We are also generating a computer plan of each building floor, so that the display can also highlight the affected room visually.

Figure 10-5 One of the alarm event display formats, showing current and previous events

10.5 Combining alarm event records with inventory data

The Department is now developing additional software in conjunction with IT Guard for use with their physical alarm system.

The current system has no knowledge of the alarmed items. The next stage of development extends the database design such that all alarmed items are barcoded. Each zone is also barcoded on its LED indicator faceplate. When an item is alarmed or moved within the alarm system, all related barcodes will be read, such that the item database is kept up-to-date. This "asset management" procedure will allow the Department to better track its equipment.

The other benefit of such a system is that the alarm display will also be able to list the items in the triggered zone. While an Electronics and Computer Science Department may generally only alarm PC's, other potential users have expressed a desire to alarm and secure such objects as paintings. By notifying security staff as to what may have been stolen, they will be more readily able to spot a potential thief leaving the building.

IT Guard believe that the inventory-linked software database system will be adopted by other academic sites, with another southern University already being a likely customer at the time of writing.

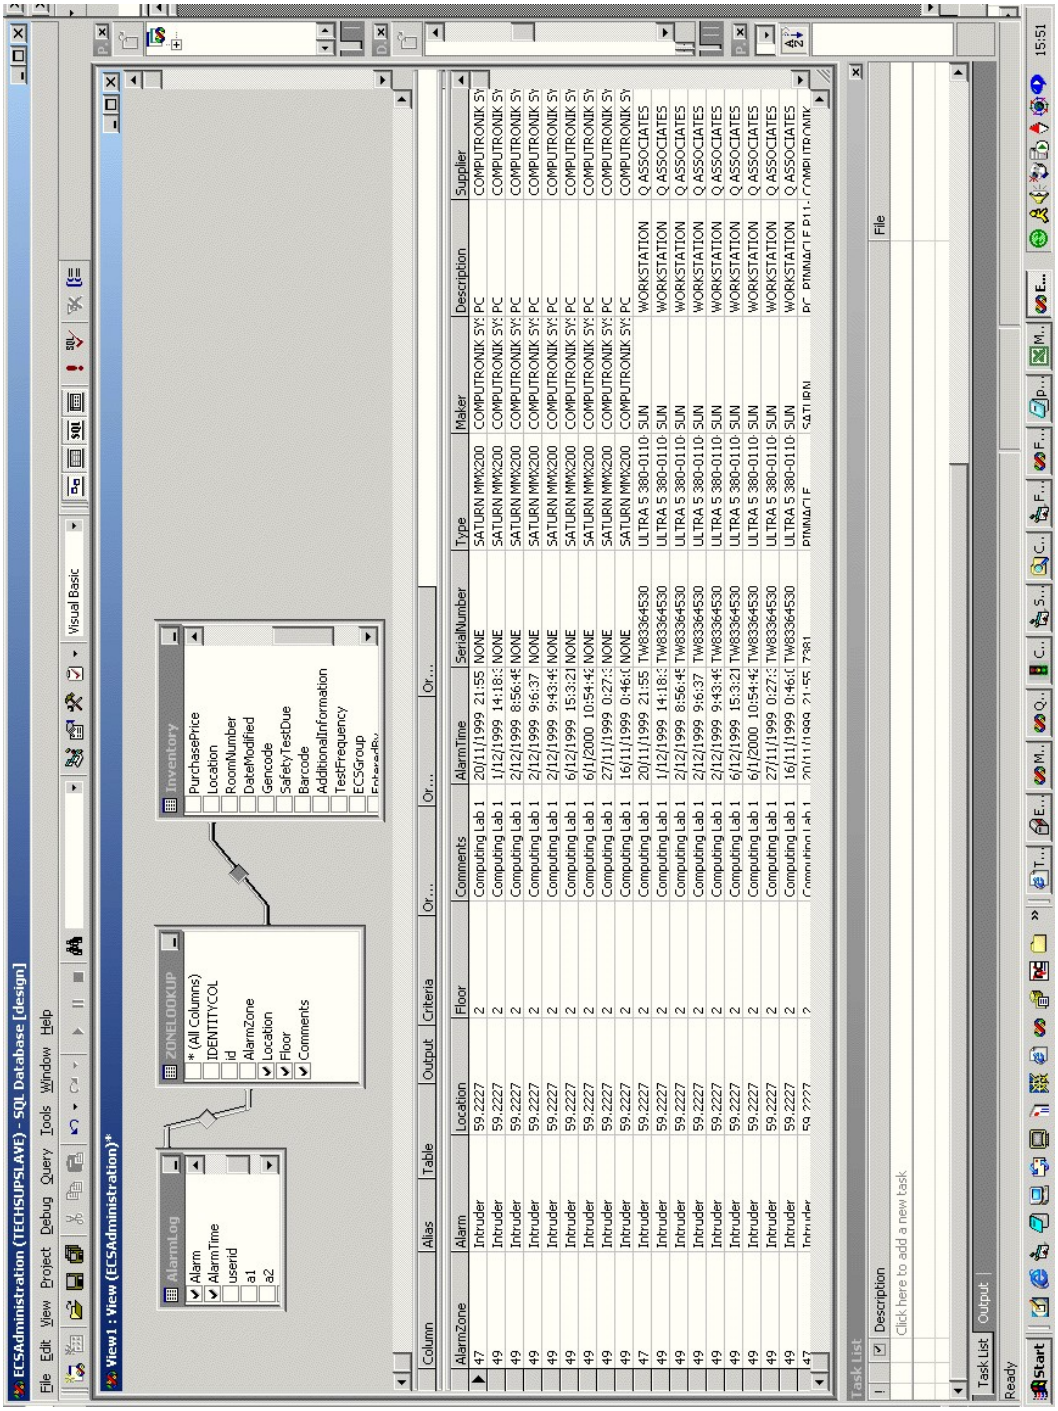

Figure 10-6 A prototype of the asset management database structures

10.6 Conclusion

A cost-effective and reliable alarm system can be run over twisted pair “data” cabling. In conjunction with the iD biscuit zone system, an accurate method for identifying a precise source of an alarm can be utilised.

Because the system runs over standard data cabling, it is possible to alarm any item by any unused data outlet by patching from the data patch rack into the alarm patch rack. In a small installation, the racks can be more readily co-located, if desired, easing flexibility.

Having developed additional software for the system we are now able to display detailed alarm event descriptions both in a public place and on a system in a security control centre, via data network protocols.

Further software developments will allow asset tracking and management of all alarmed items.

11 A holistic view of building security

The security of a building, be it in a campus environment or elsewhere, requires action to be taken on a number of fronts simultaneously. By considering multiple security aspects and addressing them as one, the overall security level can amount to more than the sum of the parts.

In considering security for the New Zepler building, the following areas were addressed:

1. *Active CCTV video surveillance*. This surveillance takes a number of forms: casually watched monitors, time-lapse VCR recordings, digital recording of "incidents" (via software such as GOTCHA!), and digital dumping of video frames from user PC's (on a volunteer basis using Webcam32 or similar software).

2. *Minimising "twilight" time by using automatic locking and card-based entry*. The card lock system ensures that staff and students can only enter the building out of hours by using their coded University ID cards. A card reader in the lift, combined with stairwell locks, restrict access to non-student laboratory floors.

3. *A single watched point of entry*. This may be through security staff outside of normal working hours, or a manned reception area during the day.

4. *Signage (and monitors) on perimeter and approach doors*. By placing clear signage that indicates that a comprehensive security system is in place, would-be thieves are (hopefully) encouraged to seek their pickings elsewhere.

Figure 11-1 Signage and card locks on building entrances

5. *Location of "attractive" equipment away from ground floors where possible*. If that is not possible, the use of blinds and/or obscuring film is suggested.

6. *The use of ID cards*. By requiring staff (and students) to wear their University ID cards on the premises, unwanted "guests" become easier to spot. Genuine visitors are assigned

temporary ID cards for the duration of their visit. With a Pan-Tilt-Zoom camera, software and a suitable printer in reception, visitor cards can be issued carrying photographs (though there may be DPA issues with such a system).

These measures do not of course address data network security; these should be addressed separately via firewalls and other measures (as to be reported by the authors' Department in JTAP-631).

We do not yet use infra-red or motion alarms due to the frequency of staff and student movements through the public entrances to our building. Contact alarms are used on important doors where it is known that no genuine visits will be required overnight.

11.1 Theft incident statistics

On the day that New Zepler was being officially opened (in mid-1999), the alarm system was tripped (audibly) three times while guests were present for the ceremony and subsequent social events. This was before we set a delay on the audible alarm trigger. On the third triggering an intruder was spotted and chased from the building. This incident serves as some evidence that an occasion where a building is open to the public can be used as some degree of cover by an opportunist thief.

The recorded instances of thefts within the Department illustrate the effect of our measures in deterring crime:

Year	Major Incidents	Value	Items lost
1995	2	£ 9,600	PC equipment, oscilloscope
1996	2	£ 6,700	4 PC's, 1 laptop
1997	3	£ 15,600	8 PC's, 1 laptop
1998	0	£ 0	-
1999	0	£ 0	-

These figures do not necessarily include all thefts, as the data used for these statistics is purely that of major claims made against insurance policies. The high excess premiums of some policies (anything from £1,000 to £2,500 per incident) mean that a claim is often not worth making.

The loss of over £30,000 pounds worth of equipment over 3 years was enough of an incentive for the Department to begin deploying anti-theft measures.

While much of the value of the stolen equipment can be reclaimed through insurance, this is dependent on the scale of the theft. A large number of single PC (or laptop) thefts will be more costly than a single incident with a larger volume of PC's taken, due to the policy excess.

Laptop PC's are a prime target, are more difficult to alarm (due to their transient nature), and often are not backed up regularly (for the same reason) so the potential cost of a loss is more than just the material cost and inconvenience.

11.2 Installation costs

The policy for physical alarm installations in the Department is to charge at material cost for the new install. The fee reflects the cost of a single sensor and Category 5 cable to connect the new item to the system, a figure not in excess of £25. The Department is thus paying for the infrastructure of the alarm system, while the Research Group (or whoever) is paying the "connection fee" (it is believed to be common for network installations to operate on the same basis).

Were an allowance for the cost of the original faceplate installation and a share of the main system included, this cost would be considerably higher. It is obviously a matter of policy for the site in question as to whether and/or how the cost of a global alarm installation is recovered.

11.3 Project cost

The cost of installing the whole security system in the Department was not insignificant. The CCTV components, not including the coaxial cabling installation, amounted to over £12,000. The card lock system was another notable five-figure sum. The IT Guard system, with cabling, ran to a large four-figure sum.

12 Related JISC/JTAP Projects

Subjects including security and video transmission over data networks have been addressed in other JISC/JTAP projects. In this section we comment on aspects of their findings and recommendations that are relevant to the work of this project.

12.1 Network Delivery of High Quality MPEG-2 Digital Video

This report includes some lengthy discussion on the problems of high quality video delivery over data networks, for teaching and learning purposes. Delivery over Ethernet and ATM is considered, as are Class of Service and Quality of Service issues. Much detail is given on MPEG and other standards, and as such readers of our report are advised to consult this report. It lacks any concise conclusions, but has recommendations such as a suggestion that Differentiated Services methods should be investigated further.

For security purposes, the problem is that with a high number of cameras multiple high quality video streams are required. Running dozens of such streams concurrently on a data network will require significant bandwidth. The report states that MPEG-2 requires 4-6Mbps for a high quality transmission, or 1-4Mbps for lower quality MPEG-2 or use of MPEG-1 instead. The report suggests multicasting as a method to allow a single stream to be fed to multiple hosts with only one instance of the transmission. However, multicast support across IP routers is variable.

The nature of CCTV footage is one where the content of the video will generally be changing little, in contrast to typical "educational" video. The bandwidth consumed by MPEG-2 CCTV may thus be typically much less than 6Mbps. There is also a decision to be made over whether CCTV is digitised at source (expensive) or on selected outputs from a video switch (for digital recording on alarm events or for feeding to a manned control centre from where the channel being viewed can be selected).

12.2 Audio-Visual Person Recognition for Security and Access Control

This report discusses biometric methods for recognising individuals, including combinations of video and speech. It includes a very useful list of current systems, and suggests a variety of applications for the technology. This of course includes authentication and access control. The report is an overview and does not appear to be based on a particular deployment case study.

One potential weakness of biometric methods which should be raised revolves around the very fact that biometrics are unique. If your fingerprint or retinal data is compromised in some way, it is not possible (yet!) to have a new eye or finger fitted. The biometric data, if not utilised and protected carefully, could potentially be replayed by an attacker.

12.3 'Videoconferencing in the Valleys': A case study of the 'ALPs' Project

This report describes a project initiated at Glamorgan whereby PC-based desktop videoconferencing was used as a means to develop "public educational opportunities for residents of the Valleys" in Wales. The project chose ISDN2 as the transmission medium, and makes no apparent reference to videoconferencing over data networks.

It is possible to hold a useful videoconferencing session over ISDN2. At the authors' site we have such a system which is used quite frequently. The advantage of running over ISDN2 is that

the very nature of a point-to-point phone call gives a guarantee of bandwidth. While this may only be 128kbits for ISDN2, in videoconferencing situations where most of the motion is restricted to lip movement, that bandwidth suffices.

The 128kbit bandwidth is generally better than the typically available cross-Internet bandwidth, though intra-JANET interconnections would be of better quality. With ISDN2 the viewer adjusts to the jerkiness of rapid movements and the lack of lip sync. The quality of the video is, however, of more questionable value for a CCTV system, but should not be forgotten as a potential (cheap) alternative to link satellite campus sites.

12.4 Overview of Watermarks, Fingerprints, and Digital Signatures

The primary focus of this report is the application of watermarking or fingerprinting with a view to protection of copyright on digital information in general. It discusses current techniques and also methods by which they can be attacked. Application of such techniques to image data is mentioned (e.g. Section 3.2c), but the report also describes those processes where the watermark is specifically designed to survive transformations to the data (as a protection of copyright), a feature not desirable in the CCTV-as-evidence domain.

The report suggests usage of a Public Key Cryptography System (PKCS) to prove the integrity of image data. As the "Digital Images as Evidence" document cited elsewhere in our report suggests, the use of a trusted third party in the data integrity policy would seem eminently sensible. Otherwise it may be argued that the means to adjust the data will have remained in the hands of the (PKCS) key holder.

12.5 Practical guidelines for teaching with videoconferencing

This report abstracts the general problem of teaching and delivery of material over a videoconferencing link from the technical aspects of the link. It does not allude to the use of either data networks or ISDN as a transmission medium (though in a recommended session checklist it does suggest you should "keep ISDN, telephone and fax numbers handy").

JANET does now run its own videoconferencing service. It also now has a Video Strategy, as published in December 1999, in which there are stated goals to

a) migrate JANET into a multi-service network
b) implement Quality of Service methods on JANET, for the benefit of video and other services.

These goals will hopefully be realised with the advent of SuperJANET4 come March 2001, along with more widespread adoption of multicast IP support.

12.6 Videoconferencing over Scotland's Metropolitan Area Network

The videoconferencing pilot on the Scottish ATM MAN, which had 21 collaborating Institutions, found that the ATM backbone provided a high quality medium for video transmissions. The report does not detail any specifics of the ATM network, in particular any use of Quality of Service methods (such as RSVP). The implied total bandwidth made available to participants is 20Mbits, though the report makes the point that the "last hop" to the desktop (i.e. over Ethernet) may have in some instances been a problem.

Throwing bandwidth at a problem is one route to a solution. Dedicated permanent virtual circuits (PVCs) offer guaranteed bandwidth, but getting bandwidth on demand (via RSVP, for example) is a more complex issue.

The free black and white CU-SeeMe reflector referred to in the report is no longer available. White Pine's current CU-SeeMe pricing structure means that the reflector (now marketed as MeetingPoint) costs a significant four-figure sum. Such a reflector, while available, was used successfully at the authors' site as a cost effective means to relay video "cues" (video of a good enough quality to spot an incident that needed investigation). There is a Linux version of the reflector available, but it is not an officially supported product.

13 Further Work – Areas of Study

This project left some of its initial goals un-met, although the focus did shift during the course of the project onto other areas.

Further areas for investigation and study include:

1. *Digital image databases*. The individual snapshots or video clips taken from digitising boards in PCs can be held on low-cost high capacity IDE drives on the capturing PC or a remote server. If stored with date and time information these can be inspected in much the same way as VCR footage. However, with the application of database and image processing techniques it is possible to post-process images or clips such only "images of interest" are held long-term, and such that the video database can be readily querid. Digital watermarking is advised in such cases, as consideration needs to be given to the continuity and authenticity of evidence requirements (as discussed earlier in this report).

2. *Quality of service (QoS) methods for video over data networks*. While low quality video (such as from CU-SeeMe) can be transmitted over Ethernet networks, transmission of broadcast quality video requires significantly more bandwidth and cannot be reliably performed over Ethernet without the use of QoS. There are standards for QoS over Ethernet (IEEE 802.1p), as there are over wider area networks (Integrated and Differentiated Service support on routers). Adoption of QoS is still in its relative infancy, and QoS between two hosts requires full compliance from all routers between the hosts. On a single site, within the scope of a single security system, this should be more easily guaranteed, but the issue of mapping Ethernet layer QoS (as seen in many Departmental data networks) to ATM layer QoS (as used on many University data backbones) has no seamless solution.

3. *Authentication.* There needs to be further study into practical "fingerprinting" techniques with general application to a variety of security applications, including digital video images or clips and other security-related data such as workstation system logs.

The advent of SuperJANET 4 (due around March 2001) offers an opportunity for UKERNA to deploy a router infrastructure where QoS is supported (whether through RSVP, DiffServ or MPLS) and where multicasting is enabled throughout. While this may not have direct application to CCTV-related video (which will liekly remain intra-campus), it is a goal that should be sought nationally.

14 Recommendations

The contents and conclusions presented elsewhere in this report lead to a number of recommendations for UK HE Instituions:

1. The 1999 Data Protection Act will affect all Institutions in a broad range of areas including both electronic and paper records. It comes into force as of March 1st, 2000. Data subjects have new rights, data controllers have new responsibilities. These apply to CCTV and video data, like any other data. Sites are urged to consult their Data Protection Officers to prepare for the implications of the new Act. One impact is that the Act will almost certainly require new procedures and policies to be documented.

2. Security of property and personnel in a building requires a "holistic" view of a suite of security measures to be taken. CCTV is just one measure. Physical alarms, ID cards, card lock access are others. All such measures should be considered.

3. A campus-wide CCTV deployment (50-80 cameras) requires significant bandwidth to be able to offer video feeds at broadcast quality. The required Quality of Service methods are not yet available, particularly to give high quality video to the desktop (or to the point of video capture). Running video over dedicated fibre is most likely the most cost-effective solution (and is one that has been adopted at Southampton).

4. A building-wide CCTV deployment (10-30 cameras) requires broadcast quality video for relays to video monitors and/or VCR. The most cost-effective method to do this is still via dedicated cabling. However, in the context of a large building or collection of adjoining buildings, it is possible to deploy a hybrid system which also utilises data networking for low quality video "cue" signals, and for digitisation of "incident" images and clips via more sophisticated WebCam software. A PC-hosted video switch, controlled over the data network, allows flexibility in the deployment and use of image capture and display devices.

5. A single camera deployment (or a small number of cameras) can be run either to a VCR or, with PC-hosted capture cards and sophisticated (but cheap) WebCam software, it can be recorded digitally. GOTCHA! appears to be an appropriate WebCam tool for such a task.

6. By setting up a secure FTP server, WebCams run on a "voluntary" basis when staff are absent from their equipment can upload potentially useful data to a safe storage location. Depending on the coverage of the camera(s), advice should first be sought on the Data Protection implications. Less sophisticated software such as Webcam32 seems well-suited.

7. Category 5(+) twisted pair "data" cabling can also be used for alarm systems and video transmission. By flood-wiring our new building with Category 5+ cabling for alarm and data use, it has been possible to wire new items into the alarm system over outlets that were originally intended for data. By over-specifying cabling requirements in a new or existing building, flexibility in deployment of extra workstations, alarm points and/or video devices can be obtained. The alarm system works on Category 3 cable, so old(er) data cabling can be re-used for an alarm system.

8. If video evidence is to be used in a court of law, sites should seek to establish audit trails and "fingerprinting"/authentication methods for their tapes and/or data. While no laws require such action, the weight of evidence presented in a UK court of law will be weak without it.

9. Security applications require accurate date and time information. Use of the Network Time Protocol (NTP) allows networked devices to maintain very accurate time. Sites should deploy NTP servers, and consider using the JANET NTP service in conjunction with any security-related computers.

10. To assist in laboratory supervision and lone working risk reduction, sites should consider the use of CCTV systems (whether traditional or electronic) to improve safety, perhaps in conjunction with "panic" alarms and/or voice-activated alarms.

11. The deployment of any CCTV system is likely to be met with protestations from some staff and students. Any deployment should thus ideally be preceded by appropriate literature and/or explanationary seminars.

12. Good practice should be followed in securing any network devices used as part of a digital surveillance system. This includes use of SSL for Web servers and secure methods for computer access and file transfer (e.g. ssh and scp). The Data Protection Act requires that reasonable precautions be taken to safeguard personal data. A digitial surveillance system could be run on a private data network, not connected to the main network (or at least firewalled from it).

13. Sites should seek to use router devices on campus which support both Quality of Service methods (e.g. RSVP and/or Diffserv) as well as IP multicast. Such support would increase the potential for satisfactory use of video transmissions over the data network.

15 Installation and contact information

All the security measures undertaken by the Department, particularly with respect to the New Zepler project, are open to inspection by staff from other HE Institutions. If you would like to arrange a visit to view and discuss the project, please contact:

Adrian Pickering
Department Manager
Electronics and Computer Science Department
University of Southampton
Highfield
Southampton
SO17 1BJ

Telephone: 023 8059 2898, E-mail: jap@ecs.soton.ac.uk

Any general queries on the report or equipment discussed in the report can be directed to:

Dr Tim Chown
Electronics and Computer Science Department
University of Southampton
Highfield
Southampton
SO17 1BJ

Telephone: 023 8059 3257, E-mail: tjc@ecs.soton.ac.uk

Information is also available on the project Web page at:

http://www.ecs.soton.ac.uk/projects/jtap633

Work on the security project has been ongoing since the JISC funding period ended. Continuous review of security processes is an important procedure to follow. On the development side, we intend to improve the event communication and handling between subsystems (e.g. the Trend card locks and the IT Guard alarms) and to continue to work with IT Guard on enhancements to their product(s).

We are again very grateful for funding by the JISC that enabled this work to be carried out to the extent that it was. We have yet to have one reported incident (theft or otherwise) since the combined measures were installed.

References

ACI International (Microswitcher PLUS), http://www.aciconnect.com

Apache SSL, http://www.apache-ssl.org/

Audio-Visual Person Recognition for Security and Access Control, JTAP Report 38, http://www.jtap.ac.uk/reports/htm/jtap-038.html, October 1999.

Axis Communications, http://www.axis.se

Brian Godette's Enhanced Reflector, http://www.dimensional.com/~bgodette/

Building Internet Firewalls, Brent Chapman and Elizabeth Zwicky, O'Reilly.

Candid Camera for Criminals, http://news.bbc.co.uk/hi/english/uk/newsid%5F191000/191692.stm, October 1998.

Dark Side of the Digital Home, http://www.zdnet.com/zdnn/special/darkside.html, 1999.

Data Protection Act 1984, http://www.hmso.gov.uk/acts/acts1984/1984035.htm, 1984.

Data Protection Act 1998, http://www.hmso.gov.uk/acts/acts1998/19980029.htm, 1998.

Data Protection Act 1998, Preparing for the New Law, http://www.dataprotection.gov.uk/prepare.htm, 1999.

Data Protection Guidance on CCTV (1984 Act), http://www.dataprotection.gov.uk/cctv_.htm

Data Protection Mailbase List, data-protection@mailbase.ac.uk

DataFellows, http://www.datafellows.com

Dimension 4 (NTP client), http://www.thinkman.com/~thinkman/dimension4/index.htm

EU Data Protection Directive (95/46/EC), http://www2.echo.lu/legal/en/dataprot/dataprot.html

GOTCHA!, http://www.gotchanow.com

House of Lords Select Committee on Science and Technology, Fifth Report on Digital Images as Evidence. http://www.parliament.the-stationery-office.co.uk/pa/ld199798/ldselect/ldsctech/064v/st0501.htm Feb 1998.

House of Lords Select Committee on Science and Technology, Eighth Report on Digital Images as Evidence http://www.parliament.the-stationery-office.co.uk/pa/ld199798/ldselect/ldsctech/121/12101.htm June 1998.

I-Spy, http://www.surveyor.com/ispy

Instant ID, http://www.instantid.com

IT Guard, http://www.itguard.co.uk

JANET Network Time Service, http://www.ja.net/ntp/

JANET Video Services, http://www.ja.net/video/

JANET Video Strategy, http://www.ja.net/development/video/strategy/strategy99_00.pdf, December 1999.

Loronix, http://www.loronix.com

Lycos Surveillance, http://www.lycos.co.uk/webguides/technology/m_spy1.html

NetCam (superVisor), http://www.netcam.ltd.uk

Network Delivery of High Quality MPEG-2 Digital Video, JTAP Report 36, http://www.jtap.ac.uk/reports/pdf/jtap-036.pdf , revised July1999.

Overview of Watermarks, Fingerprints, and Digital Signatures, JTAP Report 34, http://www.jtap.ac.uk/reports/htm/jtap-034.html, August 1999.

Practical guidelines for teaching with videoconferencing, JTAP Report 37, http://www.jtap.ac.uk/reports/htm/jtap-037.html, September 1999.

Practical Unix Security, 2nd Edition, Simon Garfinkel and Gene Spafford, O'Reilly.

Primary Image (2nd Eyes), http://www.primary-image.com

Primavision, http://www.primavision.co.uk

Privacy International, http://www.provacy.org

QoS Forum, http://www.qosforum.com

QuickCam Third-party Software, http://www.cs.duke.edu/~reynolds/quickcam/index.html

Real Networks, http://www.real.com

RS Catalogue, http://rswww.com

Security Magazine, http://www.securitymagazine.com

Statewatch, http://www.statewatch.org

Streaming Video, http://www.hwg.org/resources/?cid=38

Thawte, http://www.thawte.com

VDO, http://www.vdo.net

Verisign, http://www.verisign.com

Video Software Laboratory, http://www.video-software.com/products.htm

'Videoconferencing in the Valleys': A case study of the 'ALPs' Project, JTAP Report 35, http://www.jtap.ac.uk/reports/htm/jtap-035.html, August 1999.

Videoconferencing over Scotland's Metropolitan Area Network, JTAP Report 24, http://www.jtap.ac.uk/reports/word/jtap-024.doc, August 1998.

VNC: Virtual Network Computing, http://www.uk.research.att.com/vnc/index.html

Vortex Communications (Sourcerer), http://www.vtx.co.uk/video_7.htm

Watching Them Watching Us, http://www.spy.org.uk

WebCam Resource, http://www.webcamresource.com

WebCam World, http://www.webcamworld.com

WebCam32, http://www.webcam32.com

White Pine Software, http://www.wpine.com

Winnov, http://www.winnov.com

Winnov AV PCI card, http://www.winnov.com/products/videumavpci.htm

Winnov Videum on Linux, http://millennium.diads.com/bdirks/winnov.htm

Working Paper on Secure Internet Issues for the HE Community, http://www.jtap.ac.uk/reports/htm/jtap-032.html, June 1999

Xing Streamworks, http://www.xingtech.com/