Unit I: Introduction & evidential potential of digital devices – Key developments, Digital devices in society, Technology and culture, Comment, Closed vs. open systems, evaluating digital evidence potential. Device Handling & Examination Principles: Seizure issues, Device identification, Networked devices, Contamination, Previewing, Imaging, Continuity and hashing, Evidence locations.

1.1 INTRODUCTION & EVIDENTIAL POTENTIAL OF DIGITAL DEVICES

1.1.1 Key developments: - The time when only nerds or geeks were interested in computers is long gone. Advances in computer usability have led to the development of digital devices which are no longer the sole preserve of the white-coated "high priests" of computing (once known as the programmers and operators), but have become accessible to everyone capable of holding a mouse or using a keyboard. The IBM PC, with its standardised low-cost hardware, simple Microsoft Disc Operating System (PC-DOS or MS-DOS) and the backing of the world's largest computer manufacturer resulted in a host of imitators and compatible machines targeted mainly at business.

The success of these IBM-compatible PCs with Microsoft operating systems and applications created a de-facto standard, never before seen, which allowed free exchange of data and information between systems, people and organisations, thus eliminating one of the biggest barriers to information exchange. Standardisation of software and data created opportunities for "paper-less offices", where every member of staff had access to computing resources on the desktop – often linked to a local area network – connecting machines within an office or building for even greater resource sharing and efficiency.

Meanwhile, since the 1960s, work had been progressing on what we know today as the Internet. This wide area network began life as an academic project designed to allow data sharing between distant sites, but in a way which allowed the network to be scaled up to include millions of machines. Again, this created a de-facto international standard for networking through the creation of an easy to use system which allowed developers to add new features without compromising the existing network. In effect, the Internet provides a global "road network" which is capable of carrying any type of traffic which can be devised.

Prior to 1989, however, the Internet was largely the preserve of the technically minded, mostly because of the huge number of incompatible applications which existed on it. Tim Berners-Lee, a British physicist working at CERN, proposed a new information management system for CERN to counter the problems of information loss, damage and confusion which the organisation was suffering at the time.

The proposal defined an information sharing system which allowed disparate information systems to be linked together via a common interface based around the concept of HyperText .In a HyperText system, the user can navigate around the text by activsting links, which jump to other pieces of text. Berners-Lee's innovation was to allow these links to reference documents and even applications external to the current document. In this way, the World Wide Web as we know it was born, with a single consistent interface to a range of different applications. Arguably, this is the single most important innovation in information systems in the 20th century. It has certainly led to the widespread adoption of Internet services as a part of everyday life.

The Internet and the existence of a workable digital telecommunications network, combined with increasingly powerful low-cost devices, also created a desire to distribute more complex data in the form of music, photographs and video. Unfortunately, although the communications technologies were effective, they had not been designed with real-time high-quality audio and video in mind. As a result, it became necessary to develop compression methods

such as MPEG which would allow acceptable quality content to be delivered over low-bandwidth connections. The same technology is currently being used for broadcast digital television and is used to allow multiple digital channels to occupy the same bandwidth as a single analogue channel (although the analogue and digital signals cannot be present at the same time). Inadvertently, the MPEG2 standard for digital video had a major impact on the music industry through the creation of a new standard for digital audio – MP3 (MPEG2 layer 3).

The 1990s also saw some major changes in operating systems. Microsoft finally released its "Chicago" software, better known as Windows 95, setting a new baseline for the IBM-compatible world. This included networking in a format which was relatively easy to setup, and considerably easier than the previous Windows 3.11 and Windows for Workgroups systems, which had relied on support being provided by their underlying DOS. Developments of Windows 95 strengthened network and hardware support through Windows 98 and ME until the "home" platform converged with Microsoft's professional operating system (Windows NT) to create Windows XP and, at the start of the 21st century, Windows Vista.

Most recently, a new development in wired telecommunications has driven down the cost of high-performance internetworking to the point where it has become affordable for domestic users. Broadband xDSL technology, in the form, mainly, of ADSL (Asymmetric Digital Subscriber Line), offers a high-speed digital connection using existing telephone wiring. It offers an always-on connection, for those who want it, and allows consumers to receive more complex, "richer" content, in the form of video and other media, than was previously possible using slow dial-up connections. The increased speed also means that it has become properly possible for someone to work at home as efficiently as they could in an office. The network connection to their home computer is not as fast as the one they would have in the corporate network, but the speed is sufficient for them to access core corporate services such as e-mail.

1.2 Digital devices in society:-

The result of 40 years of innovation, as outlined above, has been a move towards an increasingly technology-dependent society. It is rare to find anyone who does not have access to some form of computing device, be it in a vehicle (engine management in a modern car), for personal entertainment (MP3 music player, CD player, DVD player etc.), personal communications (mobile telephone), personal computing or lifestyle management (Personal Digital Assitants and smartphones).

Individuals depend on digital technology to manage their personal financial affairs, ensure that goods are in shops for them to purchase and to schedule transport and other activities efficiently. We use mobile phones to communicate, wherever we happen to be, laptop/notebook computers and PDAs to work in any location and the Internet for communications, entertainment and business 24 hours a day, 7 days a week.

From a situation where activities were generally constrained to the immediate local geographic area, we have evolved into a society which operates globally but carries out activities locally. People never actually have to meet or even speak directly to each other, but can interact via e-mail, chatrooms and online ecommerce systems. Even in the preparation of this book, there has been only one face-to-face meeting. All other discussion, negotation etc. has been carried out using online systems.

Business, industry and commerce are now almost completely dependent on digital technology as a core part of their activities. Without the systems responsible for accepting and processing orders, controlling stock, issuing invoices and managing financial transactions, most businesses would start to suffer a cash-flow crisis within a matter of days.

Honest citizens and criminals alike have equal access to the technology, constrained only by cost and availability. Fortunately, most criminals make use of the technology for mainly legal purposes, but a few choose to use it to support their criminal activities. We shall explore some of the opportunities this creates in later chapters. No matter how the technology is used, however, it always records some detail of what it is doing and when it is done

1.3 Technology and culture:-

Although much technology evolution has been driven by the desire for lower power, higher capacity or greater efficiency, the emergence of consumer-oriented technologies such as Apple's iPod, mobile phones and similar personal devices has resulted in a merging between technology and fashion. In many cases, particularly among the younger members of society, it is no longer enough to have a device, but it is now necessary to have the "right" device. In the same way that people express their common interests and membership of a particular cultural group through clothing and make-up, design features of personal technology can be viewed as an expression of membership of such a group.

Indeed, since the technology has become so personal, possession of the correct device is perceived by some as an essential adjunct to participation in their chosen peer group. Teenagers, especially, seem to view their personal mobile phone as exactly that – a personal device which is guarded jealously, perhaps because it gives them a private communications channel to their peers – but only if it is the right make and model.

These personal technologies have also changed the way we communicate. There can be few people who have not witnessed the sight of a gang of teenagers walking down the street, heads down, silent apart from the hushed clicking as their thumbs fly across their mobile phone keypads sending and receiving SMS messages.

1.4 Comment:-

We seem to have arrived at a time when society is dependent on technology, not so much through need as through choice. We have driven the development of more efficient, cheaper and smaller devices because they seem to make our lives easier. The cost, however, is that we have changed the way we live to such an extent that we must have such devices in order to continue living the way we want to. Those devices, like it or not, monitor almost everything we do and can store pretty accurate records about our movements and interactions. Our technological assistants, therefore, might be viewed as unsleeping witnesses.

1.5 Closed vs. open systems

To start with, we can consider all digital devices to fall into one of two main categories: closed or open, depending on how they have been used in the past.

1.5.1 Closed systems

From the point of view of a forensic examiner, a closed system is any system which has never been connected to the Internet. This means that it has only ever existed as an isolated entity within a controlled and known environment. Any machines to which it has been connected have been closed systems themselves, thus creating a closed network; another form of closed system.

In effect, then, a closed system may consist of multiple smaller systems all of which satisfy the definition of a closed system.

1.5.2 Open systems

An open system, by contrast, is any system, no matter how large or small, which has, at some time, had some sort of connection to the Internet. This connection may have been direct (e.g. through connection to a public wireless network at a coffee shop) or indirect (e.g. through the use of a USB memory device which had previously been used in an Internet-connected system). No matter what the form of the connection and how many steps removed, any association with the Internet converts a closed system into an open system.

1.6 Evaluating digital evidence potential

The method described above allows us to evaluate the activities in which devices have participated, but it does not yet allow us to identify which devices have greatest evidential potential.

As a rule of thumb, the more categories the device falls into, the greater significance it is likely to have, particularly when those categories lie to the right hand side (Accomplice, Victim, Guardian) of the table. This is because these categories tend to represent active participants in the activity, rather than passive components.

Given the murder situation, described above, however – how would we choose which of the devices should be examined first? Obviously we would like to concentrate on the devices which contain most data of relevance, but this cannot be determined until the devices have been examined.

In this situation, we must think about the nature of the device itself, and how long data are likely to be preserved within it. A general rule, though, is that if the device is battery powered, it is probably going to lose data when the battery goes flat, so it needs to be dealt with quickly. At the very least, we need to create an image of the device so that its contents can be examined properly at some later time. No matter which devices we identify, though, it is crucial that they are correctly handled. By far the most common source of problems in digital evidence handling is the establishment and maintenance of the chain of custody.

1.7 Device Handling & Examination Principles:

Any Crime Scene Investigator, Scenes of Crime Officer, solicitor, barrister or judge will confirm that the establishment of continuity of evidence can be a crucial issue in a criminal trial. If doubt can be cast over the history of any item of evidence, allowing the suggestion that it has been tampered with during an undocumented period in its life, then the value of that item as reliable evidence is diminished. In extreme cases, it can be so compromised as to be ruled inadmissible – causing a case to collapse.

This is particularly true of digital devices as, unlike some other forensic sciences, we cannot "split" them into separate samples for testing using different processes by independent parties. The act of cutting a digital device into pieces tends to stop it working at all.

The Association of Chief Police Officers for England and Wales produces the "Good Practice Guide for Computer Based Electronic Evidence" which lays down four key principles applicable to the handling and processing of digital devices.

These principles are:

Principle 1:

> No action taken by law enforcement agencies or their agents should change data held on a computer or storage media which may subsequently be relied upon in court.

Principle 2:

In circumstances where a person finds it necessary to access original data held on a computer or on storage media, that person must be competent to do so and be able to give evidence explaining the relevance and the implications of their actions.

Principle 3:

An audit trail or other record of all processes applied to computer-based electronic evidence should be created and preserved. An independent third party should be able to examine those processes and achieve the same result.

Principle 4:

The person in charge of the investigation (the case officer) has over-all responsibility for ensuring that the law and these principles are adhered to.

Application of these principles is not restricted to laboratory-based examination, but needs to start as soon as any investigator encounters a dig-ital device. Equally, although the principles make reference to "law enforcement" and "case officer", they should be applied in any type of digital forensics work, substituting "investigative" and "manager" respectively, to give:

Principle 1:

No action taken by investigative agencies or their agents should change data held on a computer or storage edia which may subsequently be relied upon in court.

Principle 4:

The person in charge of the investigation (the manager) has over-all responsibility for ensuring that the law and these principles are adhered to.

In essence, all these principles can be summarised in four sentences as:

1. Don't modify anything.
2. If you have to risk modifying something, make sure you know what you are doing.
3. Record everything you do, in the right order.
4. Someone must take responsibility for making sure everything that is done is both legal and in accordance with these principles.

Principle 3 is particularly important in combination with Principle 2. If anything is mishandled, only accurate recording of what happened to it can allow the examiner to make allowances for incorrect handling, and perform a proper evaluation of the quality and significance of evidence recovered.

1.8 Seizure issues

It is tempting to assume that all digital devices can be seized as soon as they are identified, but this is not always the case. It would be grossly unfair, and possibly even illegal, for an investigator to seize equipment which is vital to a business as a result of suspicions of the activity of a few employees. Care must be taken to ensure that any seizure is justified, appropriate and proportionate.

Anything to be seized must, demonstrably, have the potential to contain evidence relating to the activity. It must be a major source of material and any problems presented by its seizure must be outweighed by its value in the investigation.

1.8.1 Crime scenes

At the time of seizure, it is best to consider the environment to be a crime scene and approach it using conventional crime scene procedures. Care should be taken to minimise disturbance of any items in the vicinity.

In the film versions of Richard Gordon's "Doctor" books, the surgeon Sir Lancelot Spratt was wont to say "Eyes first and most, hands next and least, and tongue not at all" about the examination of a patient. This is true of crime scenes too.

Before starting any work, it is vital that a thorough visual inspection is carried out with appropriate use of photographs and note-taking to ensure that nothing has been missed and that all risks have been fully considered.

1.8.2 Quarantine

The first step, however, has to be to establish a quarantine around the suspect equipment, moving everyone away from it to ensure that no-one has the opportunity to tamper with it. This removes the potential for any accusations of evidence being "planted" or for the user/owner to attempt to damage any evidence of which they are aware.

1.8.3 Recording status

Once the equipment has been quarantined, it should be checked to see if it is "live", i.e. showing any signs of having power applied and of software running. Its status should be recorded as completely as possible using sketches, photographs and comprehensive notes which describe exactly what can be seen. The temptation to use one's own knowledge of digital devices should be resisted. For example the phrase "A window in the centre of the screen was headed 'Microsoft Word' and contained the text 'Leave the money under the third oak tree'" may be more useful than "A word processor was running". Not every program tells the truth when it is running.

Screensavers have caused some debate amongst digital evidence practitioners in the past. Should a screensaver be stopped or simply recorded and allowed to run? Current thinking is that, since we cannot know what the "safe" way of stopping the screensaver is, it should be allowed to continue running, even if it starts during the seizure process. It is important to re-member that a screensaver is just another program in the device, triggered by a lack of user activity, and that it is capable of running other programs designed to cause damage.

Screensavers can be protected by passwords and other mechanisms. If we do not know exactly how to stop a screensaver running safely, we should treat it with suspicion and avoid doing anything to change its behaviour.

1.8.4 Networks and communications

It may be obvious that the machine is connected to a communications device such as a network port or modem. In this case, opinion is divided about the best action to take. General advice, though, is that in the absence of specialist assistance, the safest option is usually to disconnect the system from communications devices as quickly as possible, often before the status recording is complete. However, disconnecting a system from a live network/communications session presents some risks.

Firstly, if the system is connected to a co-conspirator in some way, the disconnection of the communications without warning may alert members of the gang to the fact that something untoward has happened. This may give them time to destroy evidence of their own involvement before they have been identified.

Secondly, it is possible for any system to detect loss of communications and commence action ranging from deletion of data to, in theory, triggering an explosive device.

Finally, in the case of mobile phones, switching the phone off to remove it from the network causes the phone to change internal data which might have been useful to the investigation. It is better, when at all possible, to seek advice from a specialist about how to deal with live communications, but if this is not possible, accurate recording of the actions taken should allow the laboratory-based digital evidence examiner to account for any anomalies caused by actions taken in the field.

1.8.5 Power

Having conducted and recorded a thorough visual inspection, and dealt with the issue of communications, we can turn to the issue of shutting equipment down. It is recommended that all systems to be seized should be shut down as soon as possible after their discovery.

Using the operating system's (O/S) "shut down" or "halt" command is not a good idea, nor is simply pressing the power button.

We are all taught, almost as soon as we start to use computers, that it is vital to shut them down correctly. In most cases, this means using special commands to ensure that all data have been written to storage devices correctly. Anyone who has ever removed a floppy disc from a drive before the light has gone out, or taken a USB memory device out of a Windows PC without using the "safe remove" option will be aware of just how bad the damage can be. Notwithstanding the potential for damaging storage devices, it is still far more dangerous to shut systems down "cleanly" during seizure for two principle reasons.

Firstly, and perhaps most importantly, the O/S shut down is a software process, potentially made up of several programs. Each one of these may cause data to be written to the storage devices. As soon as this happens, ACPO's first principle has been violated and we have created a problem with the integrity of the devices.

Secondly, the shut-down process may not be the same on all machines. Because shut down is a software process, it can be modified at will by knowledgeable users. They may, if they choose to, plant programs in the process in order to damage evidence – or worse.

So, why not press the power-off button? In modern systems the power button is actually used as a trigger for the O/S shut-down process. Even though pressing it for 10 seconds (or slightly longer) seems to kill the system completely, during that 10 second period, software may start to run and cause damage.

By far the preferred method is to remove power directly from the system by disconnecting it from the mains. Even this is not as straight forward as it may sound.

Instinctively, the correct way to do this would appear to be by switching power off at the wall socket and then removing the plug.

Unfortunately, doing this simple action may send a signal to the device that power has been lost and allow it to take action to damage potential evidence. This is because of the presence of an Uninterruptible Power Supply (UPS) (see Figure 3.1).

Figure 1.8.1 A typical uninterruptible power supply

A UPS is a battery which is kept charged from mains power and which is designed to supply power to a device even when mains power has been lost. Typically, it is used to allow servers and other systems to shut down cleanly when power is lost, thus minimising the chances of damage occurring. In order to allow the server to take appropriate action, the UPS has to be able to send it a signal informing it that a power-cut has occurred. This signal triggers software activity, in exactly the same way as pressing the power button, or running the shut-down process.

For these reasons, then, the recommended process is to pull the power lead from the socket on the device itself (Figure 3.2), or as close as possible to the device. This will ensure that there are no external UPS or similar power sources in circuit. Of course, there still remains the problem of internal power sources.

Note, however, that accepted practice is to allow any printing or CD/DVD writing to finish before disconnecting the power, as these produce a near-permanent record of activity that was happening when the machine was discovered.

Once the device has been isolated from the power supply, all leads and their associated sockets should be clearly labelled (see Figure 3.3), in case it becomes necessary to reconnect devices during laboratory examination, and everything should then be packaged and labelled to start the evidence audit/continuity trail (see Figures 3.4 and 3.5).

From this point on, whenever the device is handed over to any person, the label should be completed showing the date and time of handover along with the identity of the person it is passed to.

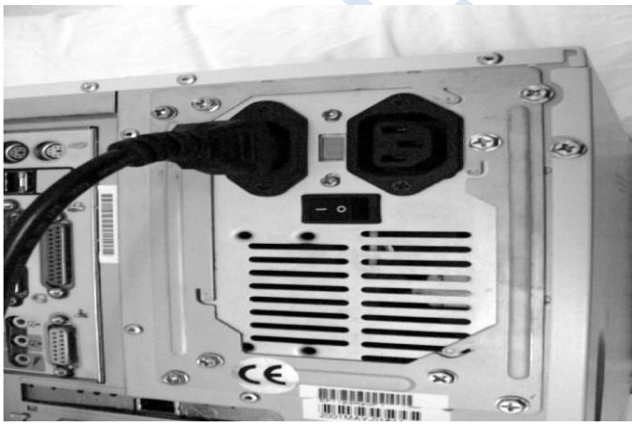

Figure 1.8.2 A standard power connector on the rear of a PC

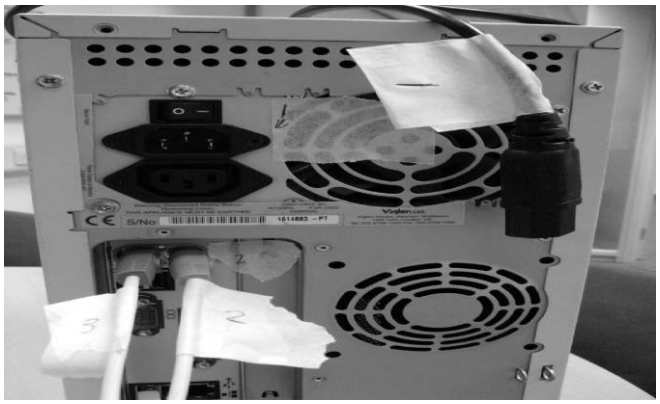

Figure 1.8.2 Labelling the sockets and cables before disconnecting and packaging

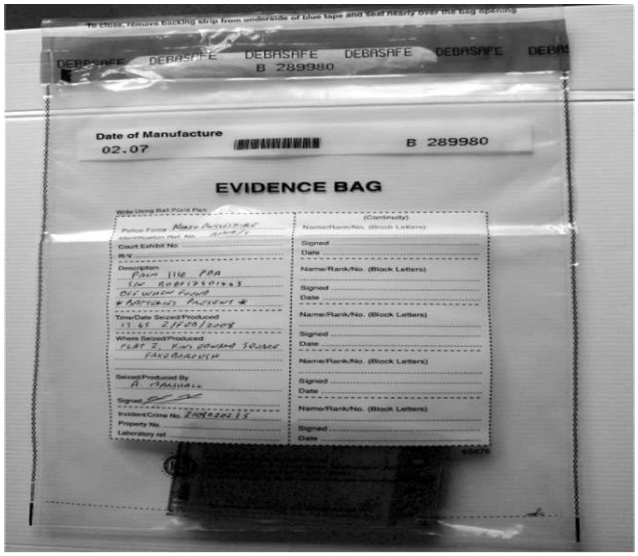

Figure 1.8.3 A PDA in tamper-evident packaging

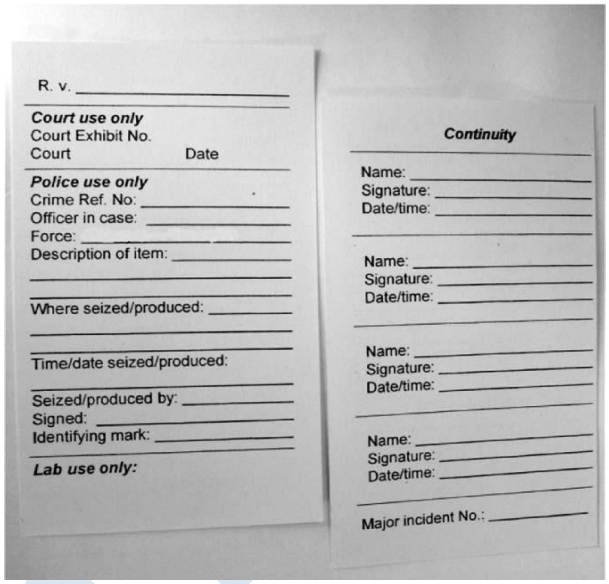

Figure 1.8.4 Front and rear of a sample evidence label, also know as a "CJA" (Criminal Justice Act) label

1.8.6 Threats and risks

Throughout the description of the seizure process, above, there has been constant mention of the fact that almost any interaction with a device being seized can cause changes to the state of that device. This is a real risk. If the state of the system, when it comes to be examined, can be shown to have changed during or after seizure, then the integrity of all data on that device can be challenged, effectively accusing someone involved in handling the device of tampering with it.

1.9 Device identification

From the discussion above, it may sound as if it is easy to identify digital devices. Most people have a very clear mental picture of a PC as a beige or black box with keyboard, mouse and monitor attached to it. Similarly mobile phones, media players etc. seem to be easily identifiable devices. However, the truth is that the onward march of Moore's law means that the "chips" at the heart of all devices are getting smaller and more powerful all the time. As a result, more features can be packed into smaller packages.

Combine this with the fact that digital devices have become lifestyle devices rather than purely technical solutions, and we have to incorporate the influences of design and fashion trends into any consideration of device identification.

1.9.1 Case modding

Amongst a certain group of owners, there is a fashion for case modification, aka "modding". Much as owners of mundane vehicles attempt to make them more attractive or unique by attaching spoilers, neon lights, chrome wheels and fancy paint schemes, case modders change the external appearance of their computers to make them more personal or attractive.

Web sites such as http://www.mini-itx.com/ (see Figure 3.6) contain de-tails of readers' projects showing how they incorporate PCs into everything from whisky bottles to dolls based on Japanese manga characters.

As a result, PCs may no longer look like PCs, but can be disguised as household objects. Indeed, as the convergence between entertainment equipment and computing equipment moves forward with O/S such as Windows Media Center Edition allowing PCs to be used as video recorders and hi-fis, there is a desire to "improve" the appearance of computing equipment to make it fit better within the living environment.

Another approach taken to the problem of making computers more suitable for domestic use is that taken by manufacturers such as Shuttle and Apple, both of whom have produced "mini" computers which occupy little more space than a small stack of CDs and which are designed to look like ornaments or conventional small hi-fi equipment.

The only common factor with these devices is the need for them to have cables for power, video, network, keyboard and mouse. The last three of these requirements are disappearing, too, as wireless keyboards and mice become more reliable and wireless network speeds and reliability increase.

Figure 1.9.1 Sample from http://www.mini-itx.com/ showing a model aircraft carrier which contains a full PC

1.9.2 Novelty items

Another category of device which can pose problems is that of portable storage. Typically, these are lumped together under the catch-all terms "thumb drives" or "USB sticks" because the earliest and most common versions of these USB devices were about the same size and shape as a thumb (see Figure 3.7).

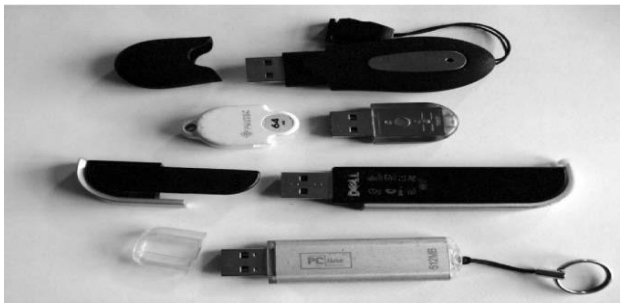

Figure 1.9.2"Standard" USB storage devices

Moore's law plays a part here again. As the chips for these devices have shrunk in size and/or increased in capacity, the real limiting factor in their external design has become the USB connector itself. Outside of the requirement to have the right physical connector, there are no limits to the shape, size, colour etc. of these devices. Everything from wristwatches to toy dolls, via pens, fish-fingers, sushi and keyring footballs has been used as a casing for solid-state USB storage (see Figure 3.8).

1.9.3 Purloined letters

Although it is becoming easier to disguise devices through the production of customised cases, it is still easier to adopt Edgar Allan Poe's "purloined letter" approach where the object is, effectively, hidden in plain view but disguised by being placed in such an obvious location as to appear not hidden, or is disguised as an inoccuous object.

As an example, consider a DVD on which a collection of illegal images of children (IIOC) has been written. The criminal may take steps to hide this under floorboards, behind a panel in the wall, or in some other secret location. Alternatively, he may choose to hide the DVD by simply placing it in a case for a commercially available DVD, possibly even going as far as printing a false label on it to further disguise its identity. If the illegal disc is then placed in a collection of innocent discs, its presence will be less obvious and more likely to be overlooked if a search of the disc collection is anything less than thorough.

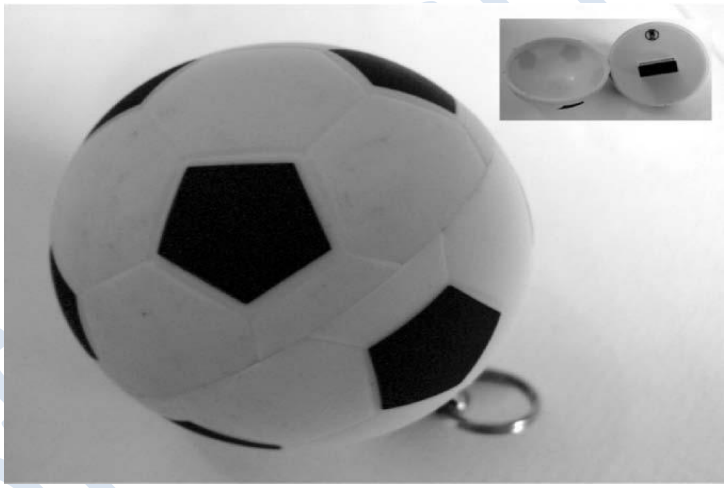

Figure 1.9.3 A novelty/disguised USB storage device (inset shows the USB connector visible when it is opened)

This method has been employed to hide several different types of device and, with miniaturisation leading to devices such as micro-SD or Trans-Flash cards (see Figure 3.9) it has become possible to secrete large volumes of data in very small hiding places such as the spines of hardback books, children's coin banks and so on.

A further twist lies in the fact that storage devices can be shared by several other devices so that a single SD card (Figure 3.9) may contain photographs which are visible when it is used in a digital camera, but the same SD card may contain other images which are not accessible by the camera and only available when the card is used with a mobile phone or computer with card reader.

The underlying message, here, is that objects are not always what they appear to be and should be checked, thoroughly, to see if they may contain any digital devices. Equally, digital storage devices should not be assumed to contain only one type of data simply because of where they were found.

Figure 1.9.4 SD card (large) and micro-SD/TransFlash card (small) devices

1.10 Networked devices

As mentioned earlier, networks can pose a particular problem during search and seizure. any form of communications immediately introduces the possibility that we are dealing with multiple machines. Where a wired network is present, it may be sufficient to trace the cables from the network socket (Figure 3.10) to the router (Figure 3.11) or hub/switch (Figure 3.12) used to connect the various machines together.

Cable tracing, though, may be complicated by the fact that cables can be run under carpets, below floors, through walls and through ceilings. Each time a cable passes through a wall, floor or ceiling its direction and point of exit may not be clear. Indeed, there may be a hidden hub, switch or router used to extend the network and split the connection to lead to even more machines. Irrespective of this, cable tracing should always be attempted. In this process, the presence of any of the devices mentioned can be helpful as the lights on the front of these devices give a good indication of which cables are live and carrying data. If all cables have been disconnected from all but one obvious device but the switch is still showing data being transferred, there must be at least one more device attached to the network.

Figure 1.9.5 A network socket on a PC and associated cable

Figure 1.9.6 A domestic broadband router with wireless antenna

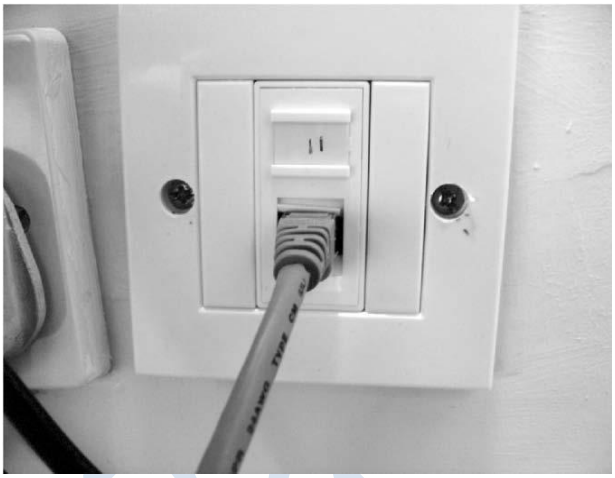

Figure 1.9.7 A low-cost network switch used to connect multiple machines **Figure 1.9.8** A standard RJ45 wall plate in the author's bedroom

In the event that the network has been installed using standard RJ45 wall sockets (Figure 3.13), it is likely that all the wiring in the premises leads back to a central patch-panel (Figure 3.14) or switch (Figure 3.12) from which all live circuits can be located. This is most commonly found in business premises, but this method of network installation is becoming more popular in domestic situations.

1.10.1 Wireless

A relatively recent innovation is the introduction of the two common wireless network protocols known as WiFi and Bluetooth.

Figure 1.10 A network patch-panel in a domestic installation

The IEEE802.11 family of wireless networking protocols for local area networks, often referred to as "WiFi", defines a set of standards using different carrier frequencies for medium to high speed networking over a limited range of 100 m (or more) in clear air (i.e. outdoors in ideal conditions). It offers similar features to wired networks, although speeds tend to be lower, without the requirement to physically plug-in to a network point. As such, it has become very popular in domestic settings as a way of sharing broadband connections without having to run cables to every room, and as an added service for travellers in cafes, coffee shops, hotels, airports and other places where it may be useful to use an ad-hoc network connection.

Bluetooth, meanwhile, offers a "personal" area network (PAN) de-signed to allow items of personal technology to work at short range (typically 1 m to 10 m, although the standard does allow for up to 100 m). The most popular use of Bluetooth, currently, is for hands-free headsets for mobile phones, although printers and various other devices are available in Bluetooth enabled form. If either of these technologies are in use in the vicinity of suspect equipment, it may be useful to know this. Since both are based on broadcast radio frequencies, it is possible to detect them using appropriate receiving equipment, such as a notebook/laptop computer equipped with an appropriate transceiver card and specialist software which scans for the appropriate frequencies without attempting to connect to detected networks. Tools such as NetStumbler and Kismet have proved useful for the detection of IEEE802.11 networks, while BlueStumbler and similar tools perform the same function for Bluetooth systems.

1.10.2 Remote access

Whatever the type of network present, there exists the possibility that some-one outside the premises has access to the equipment under investigation through the Internet. The ACPO Good Practice Guide [2] recommends that connections to the Internet should be disconnected immediately upon discovery.

Where no other advice is available, the ACPO recommendation is appropriate; however it should be remembered that a sudden disconnection from the network may alert remote users to the fact that the machines are under investigation. In an ideal world, monitoring of live Internet connections should be considered, although there is some debate about whether this constitutes wire-tapping or some other form of covert surveillance, with consequent concerns about admissibility of evidence gathered.

1.11 Contamination

In "conventional" forensic science, one of the major concerns associated with the handling and examination of any item of evidence is that of contamination or cross-contamination. The possibility of fibres accidentally being transferred between items, DNA mixtures appearing or fingerprints being deposited after an item has been seized have all caused problems when evidence has come to court in the past.

In the field of digital evidence, similar concerns about contamination can arise and can be considered in two distinct areas: physical contamination and digital contamination.

1.11.1 Physical contamination

Although it is still not common practice for digital devices to be examined for physical evidence such as fingerprints, fibres and DNA, some work in this area has suggested that physical evidence may have a role to play in corroborating or disproving theories about the physical processes under-gone by devices. For example, examination of fingerprints and tool marks on a hard disc may assist in determining if it has been replaced by the owner or if it is still the original installed at the factory. Consideration of the fibre population inside a PC may give some indication of whether it has been moved recently, and examination of fingerprints and earprints on a mobile phone may help to identify the most recent user. For these reasons, care should be taken to ensure that opportunities for physical contamination are minimised during packaging, transport and laboratory examination of devices.

1.11.2 Digital contamination

Digital contamination is another serious issue. ACPO Principles 1 and 2 require that steps must be taken to ensure that material held on storage devices is not modified in any way unless absolutely necessary. Unfortunately, the presence of wireless network equipment, with coverage that extends beyond the physical boundaries of the premises containing the equipment, presents new challenges.

It has become almost impossible to obtain notebook computers without WiFi capabilities, or mobile phones without Bluetooth. Technologies such as these present several opportunities for damage to digital evidence, either deliberately or accidentally.

A knowledgeable criminal may deliberately choose to use the technologies as "sniffers" constantly seeking new devices and, when a new device is detected, trigger software which destroys evidence held on the machine before it can be seized. Worse yet, a terrorist could use the same technique to count the number of unrecognised wireless devices within range and use this information to trigger an explosive device only when the threat to human life is maximised.

Another approach is to "pair" pieces of equipment using wireless methods. If one member of the pair is seized and taken out of range of the other (or even switched off), this can be treated as an unexpected event and, again, used as the trigger for some activity.

In all of these cases, the aim of the criminal is to gain advance notice of unwanted attention and to take steps to destroy incriminating material. The use of the wireless technology allows the process to be automated.

Another risk, though, arises from the possession of wireless-capable devices by those involved in investigating crime and seizing equipment. If care is not taken to ensure that the investigators' devices have had their wireless functions disabled prior to approaching a suspect device, there is the risk, not only that their devices will be detected as above, but also that their devices may carry out default behaviours to join local wireless networks. If the investigators' notebook computer, for example, is allowed to join the criminal's computer network, the state of machines on the network will change and ACPO Principle 1 will be violated. Similarly, a Bluetooth-enabled phone may attempt to pair with any other phone, headset or other Bluetooth device as soon as it detects its presence.

Ideally, then, no wireless-capable devices will be permitted anywhere that digital devices are thought to exist, unless they are in the hands of properly qualified personnel who have good reason to use them (e.g. in order to detect the presence of other wireless devices).

1.12 Previewing

The preview process may be considered to carry a significant risk of violating the ACPO's first principle, because it requires direct examination of the suspect devices. However, correct application of Principles 2, 3 and 4 pro-vides some degree of protection from accusations of evidence-tampering.

1.12.1 Offline preview

In a typical preview situation, the device to be previewed will be dealt with in an offline state. That is to say, it will have been disconnected from networks and shut down to allow the examiner to remove storage devices for connection to a trusted preview workstation. Ideally, devices will be connected through a write-blocking device to ensure that ACPO Principle 1 is still upheld.

This used to require specialist hardware for each device interface type in common use. A more modern approach, though, relies on the existence of the USB mass storage standard, which allows many different devices to be connected through the same standard USB interface. Because all these devices use the same protocol for communication, it is possible to use a common USB write blocker to protect them against accidental digital contamination.

Figure 1.12.1 A Tableau T8 USB write-blocker used to protect devices against accidental data writing

Figure 1.12.2A typical SATA/IDE-to-USB hard disc adapter

The write-blocker is connected directly to the forensic workstation and the device under examination is then connected, possibly through an adapter such as that shown in Figure 4.2, to allow previewing to proceed.

Once the device has been correctly connected, previewing is conducted using software similar to that used in the laboratory, to gather information about the data held on the device and allow the examiner to judge the likely evidential value of the device as a whole.

1.12.2 Online preview

The benefit of offline previewing is the opportunity to use physical methods to prevent contamination of possible evidence. The major disadvantage with it, however, is the requirement to shut down the device under examination and gain physical access to the storage devices it contains.

In some situations, it is not possible to shut down the equipment and we need to consider an alternative approach: online previewing.

Conducting an online preview is, arguably, the most risky activity any digital evidence examiner can perform. By its very nature, the preview examination will be carried out on a live system, which is undergoing changes of state and which cannot be considered to be completely trustworthy.

Although some degree of previewing can be conducted using programs already present on the suspect system, it is difficult to show that they are giving complete and accurate results. Programs such as rootkits, which are designed to implant themselves into a system to allow an attacker to abuse that system, tend to contain features which alter system programs and functions to disguise the presence of the rootkit. Thus, any program which makes use of system functions must be considered to be inaccurate unless it can be proven otherwise.

In order to conduct an online preview, therefore, most examiners have access to trusted tools on read-only media such as CDs. These tools have been written in a trusted environment and the CD contains all the code necessary to allow them to run without using any code from the suspect system. They tend, also, to be written in such a way that they protect the suspect system from accidental changes to data, apart from the necessary changes to primary RAM[1] to run the trusted tools themselves.

The aim of previewing is, generally, to establish something akin to "probable cause" in order to determine whether or not the equipment being previewed can legitimately be seized for laboratory-based examination. Although every attempt is made to follow the ACPO principles, it can be difficult to prove that the previewing process has not caused changes to data on the system and it may be difficult to use evidence acquired in this way for anything other than intelligence purposes.

1.13 Imaging

As suggested above, the usual recommendation is that any forensic examination of a system should concentrate on storage devices, and that all work should be carried out on copies or images of those storage devices. This is similar to the requirements of the HOSDB Digital Imaging Procedure, which mandate that before any work is carried out on a digital photograph a write-protected master copy should be created and preserved and further working copies generated from it when necessary.

To produce an accurate copy of a digital storage device, we need to use a method which will produce a complete copy of the device, including all unused space, deleted data and, if possible, damaged areas. This is not a normal function of any standard software available to most users. Instead, specialist tools or standard tools with special options are used to get the most complete copy possible.

1.13.1 Offline

Offline imaging is the simplest procedure, although it can be time consuming depending on the size of the device to be imaged. In this process, the suspect device is connected to an imaging workstation using a write-blocker, The imaging software is then used to read data from the device and store it to either a file or separate device. Once imaging is complete, the first copy is usually considered to be the master copy and further working copies can be generated as required. Of course, there is still a requirement to show that the master copy and working copies are completely accurate and have not been modified in any way during imaging or examination.

1.13.2 Online

When it is not possible to seize or shut down the suspect system, or where the storage device is difficult to connect to the imaging workstation, it may be necessary to use online imaging. In this situation, the storage device is left in situ and live imaging tools are used to capture data from it using the accompanying hardware. Of course, this poses similar problems to online previewing, but the use of trusted tools goes some way to mitigating these. The trusted tools allow the examiner to copy the device to either an external storage device, such as a USB hard disc, or across a network to a dedicated storage server.

1.13.3 Backups

One final method can be used to collect data from a suspect system and is particularly applicable in business environments. Backup tapes or discs produced over a period of time can allow the examiner to build up "snap-shots" of system states at the time backups were produced. Although these will contain only files which were live and scheduled for backup, the mate-rial contained in them has proven useful in the past. It must be remembered, though, that any image produced from backups is only a partial image and that hidden data is unlikely to be found in this situation unless multiple backups can be used to produce images of the system at various points in time.

1.13.4 Continuity and hashing

Once an image of a storage device has been created, the image and the device need to be treated as if they were crime scenes, with all that means for preservation of evidence and avoidance of contamination and tampering. In the

physical world, this is achieved by establishing cordons/quarantine zones and ensuring that all activities are subject to thorough recording.

The use of write-blocking methods should ensure that original devices are not subject to modification, but additional checks can be used to demonstrate that the processes used by the examiner have had no adverse effects on either original or image. These methods can also be used to alert the examiner to any accidental changes, allowing him/her to check and recheck processes which seem to be problematic.

1.1 EVIDENCE LOCATIONS

Table 1.1 Hash values for strings which differ by one bit

String	*Computed MD5 Hash (in Hexadecimal)*
12345678	23cdc18507b52418db7740cbb5543e54
12345679	0f4fd7804fbbcf67df5dc8ef8dc946fb
22345678	0c7e888e4e214b74c1ec2b6734096fe6

Hashing algorithms similar to checksums are used to calculate digital "signatures" which are effectively unique for any piece of digital data. At the highest level, one or more hash values will be computed for the data on the original device. Because the image has been produced from this device, and contains identical data, the hash value for the image should match, exactly, the value for the original. Hashing algoriths such as MD5, SHA-1 and SNEFRU are very sensitive to changes in data, and the modification of even a single bit (1/8 of a byte) in the largest image results in radically different hash values being calculated.

Some tools use hash values for sub elements of an image to allow the location of any modification to be found more easily, and provide additional assurance of evidence integrity.

Good tools also provide comprehensive logging facilities, allowing them to produce a complete list of all actions performed on the evidence during examination, helping to fulfill the requirements of ACPO Principle 3.

1.14 Evidence locations

In examining any device to recover and analyse material present, there are four primary types of file or data which the examiner will typically consider. These can be summarised as:

1. Live data
2. Deleted data
3. Swap space
4. Slack space.

1.14.1 Live data

Live data are the data present on a system in a format which makes them accessible to the user or the normal software directly. Typically, they represent the outcome of some normal operation of the device or software as a result of deliberate action.

Generally speaking, live data have greater evidential value as they can be shown to be directly related to something the user of the system has chosen to do. Furthermore, because live data files are created and managed by

the operating system on behalf of the application software, they tend to have reliable timestamps, insofar as the device's clock can be trusted.

1.14.1.1 Timestamps

Most operating systems maintain three timestamps for each file on the system, known as the *MAC* (Modified, Accessed, Created) times:

1. Modified: this records the time that the file was last modified, i.e. when it was last saved.
2. Accessed: records when the file was last read. On many operating systems this only records the date of reading, not the time on that date.
3. Created: records when the file first appeared on the system as a new file.

These timestamps are just one area of file *meta-data* which can prove useful in determing the sequence of events and nature of activity on a system.

1.14.2 Deleted data

Following on from live data, perhaps the next most useful area to consider is the deleted data on the system. This represents material which was live at some point in the past, but which the user or operating system has chosen to remove from the system for some reason.

Fortunately for the forensic examiner, most operating systems do not completely erase deleted data. Instead, they typically mark the area of the storage device occupied by those data as available for re-use. This is done because, historically, deleting data was a large, time-consuming task and it was (and still is) far easier to simply mark space as available for re-use and overwrite it when necessary.

The disadvantage of this, for the criminal, is that there is a good probabil-ity that old data can be recovered using appropriate tools. Unfortunately, because the files have been marked as deleted, their meta-data can no longer be considered entirely reliable, so information such as the MAC times cannot be relied upon in the same way as it is for live data.

1.14.3 Swap space

Another common feature of modern operating systems and applications is the ability to use storage devices as if they were part of the machine's primary RAM. Historically, and today, this has been done to make a machine appear to have more main memory than is physically present, as a way of keeping the cost of hardware down. Real primary RAM costs, on average, about 10 times as much as disc storage.

The method is similar to that of dealing with a phone call whilst in the middle of attempting to solve a difficult problem. As soon as the phone starts to ring, signalling that a new task is about to begin, the smart thing to do is to write down information about the problem and thoughts so far. Once the phone call has been dealt with, ideas can be reloaded into the brain by reading the notes made before the phone was answered.

The swapping process is completely automated and under the control of software. The user can usually only choose how much disc space can be used as swap space, but cannot choose which programs or pieces of data will be swapped out and in.

As more data and programs occupy main memory, the system monitors memory usage and identifies those areas which are least used in terms of frequency of use, recency of use or some other measure. When main memory becomes too

"crowded" to accomodate new software or data, the operating system chooses which parts of memory to store to the disc and re-allocate for new data or programs.

Later, when the program which originally owned the memory needs access to it again, the operating system chooses a new area to swap out, stores this to disc and then loads the original data back into main memory.

Since the user has no control of the process, it is possible to find data in swap space which has never been written to storage devices for any other reason. Swap space is often a good source of information about passwords for encrypted files and other sensitive data.

Because swap space is in constant use, and has no particular meta-data associated with it, it is difficult to determine exactly when data were deposited within it.

1.14.4 Slack space

Finally, we can consider slack space: data which are stored alongside live data, but not deliberately put there by the user.

Slack space tends to arise because of some properties of storage devices and file systems which dictate that there is a minimum quantity of data which must be written to, or read from, a storage device. The exact size of this Minimum Allocation Unit or MAU varies from file system to file system. The larger the MAU, the greater the chance that slack space may contain useful data.

Figure 4.3 shows how the use of MAUs leads to slack space being created when data are written to the sorage device. If the program requesting the write provides insufficient data, the system is free to grab data from anywhere in memory and use it to pad the data to be written to ensure that the MAU is completely filled. Again, the user has no control over how the system chooses which areas of memory to use as padding.

In this case (Figure 4.3), three blocks are completely filled, but there is slack at the end of Block 4, which must be filled with other data from some location in RAM.

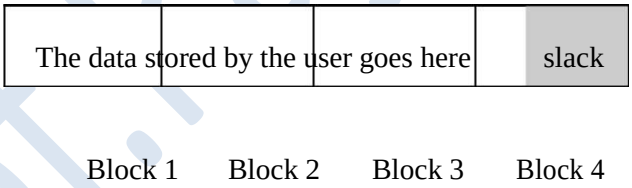

Figure 4.3 Slack space

Because data in slack space are written at the same time as the real contents of the file, the modified time in particular can be used to show that the data were present in main memory at the time the file was written. The data in slack space are not made available to programs as a result of file reading, so the accessed time has little or no significance in this situation.

Unit II: A seven element security model, A developmental model of digital systems, Knowing, Unknowing, Audit and logs, Data content, Data context. Internet & Mobile Devices, The ISO / OSI model, the internet protocol suite, DNS, Internet applications, Mobile phone PDAs, GPS, Other personal technology.

2.1 A seven element security model

This model is based on a holistic approach to security, acknowledging that security depends on all elements of the system being secure in their own right, and showing the inter-relationships between key parts of the system. It contains seven key elements which are inter-linked to compose the system as a whole. Each of these elements must be secure in its own right and depends on its neighbouring elements to ensure this. If there is a weakness in any one of the elements, the whole system is insecure and may be attacked/abused through the weak element(s).

2.1.1 Entities

Entities are objects which can manipulate and/or be manipulated by the system, and may be passive or active. Passive entities are, generally speaking, external to the system, and are represented by collections of data which the system processes in some way. Active entities tend to be part of the system and are responsible for choosing and controlling which processing should be performed. Entities can be people, organisations or other types of object.

In security terms, we need to consider the behaviours of active entities and ensure that they are only allowed to behave in inherently secure ways. We also need to think about how changes to the system affect entities and whether those changes have been properly assessed for risk. Finally, the inter-relationships between entities may lead to insecurity through inappropriate interactions or through inferences being drawn about unknown entities based on information provided about known entities.

From an investigative point of view, entities are generally either the users of the system, some of whom may be involved in the unwanted activity, or they are potential targets of the unwanted activity because they have some value to the criminal.

2.1.2 Environment

In this context, the environment represents the set of constraints or restrictions imposed on entities in an attempt to make them behave correctly. Depending on the entities under consideration, the environment may be defined in terms of: legislation, ethics and regulation; technical capabilities and resource limitations; compatibility with other entities, interaction standards and procedures; or physical limitations. Examination of the constraints and restrictions contained in the environment may allow us to quickly identify any unconsidered vectors of abuse present in the system and, indeed, determine if it is feasible to proceed with an investigation.

In some circumstances the environment may prevent us carrying out an investigation because the activities required would break some of the laws which make up that environment. Equally, if we cannot identify applicable laws, policies or rules which prohibit the activity under consideration then, no matter how

thorough the investigation is, there can be no action taken against the perpetrators. We can, of course, take action to ensure that the unwanted activity does not happen again.

2.1.3 Organisation

While the environment represents the complete collection of restraints and controls, the organisation contains the framework which allows those restraints and controls to be created, enforced and inter-related to each other while still allowing entities to co-operate and collaborate as necessary. If the organisation is inadequate, then any investigation will struggle for lack of support, co-operation and resources.

2.1.4 Infrastructure

Infrastructure is the supporting mechanism required to enable activities within the organisation. It is generally composed of physical components (such as buildings, power supplies, computers, network switches etc.) which do not need to be considered as entities, but rather as enabling hardware. In some environments, the infrastructure may be more nebulous, for example the trading floor of a stock exchange would be an infrastructure, but it still has its own security requirements which act to protect the elements identified.

When considering an investigation, the infrastructure, or some part of it, is the crime scene that will contain the information we are seeking. In order to be able to identify the limits of the scene, we must have access to an adequate "map" of the infrastructure, showing how it is structured, in order that we can attempt to identify the boundaries of the critical components.

2.1.5 Activities

Activities are complete tasks, often representing a complete transaction, from end to end. Activities can be single procedures, but may be composites made up of multiple procedures. Insecure activities are those which are poorly defined or understood and which have the potential to allow unwanted side-effects to occur. A good understanding of the activities in progress during the incident being considered should allow the investigator to reconstruct the sequence of events which led up to the incident, taking account of unexpected behaviours.

2.1.6 Procedures

At the procedure level, we are considering single tasks which represent discrete, identifiable, atomic (indivisible) processes within the wider context of an activity. For example, making a cup of tea is an activity, with a clearly defined start and end point, which is made up of procedures such as "putting water in the kettle", "putting tea in the pot" and so on.

Procedures should be the smallest steps involved in an activity which may be shared with other activities (e.g. "putting water in the kettle" is shared with "making a cup of coffee").

Really, here, we are applying classical top-down design methodology to understand more about the actions taken within the system. By breaking them down into procedures which can be described simply and quickly, we increase the chances of identifying any that are weakly defined or more likely to result in unwanted or unexpected results.

2.1.7 Data

Finally, we come to data: the representation of the entities which the system is concerned with manipulating. Two key areas of security apply to data. Firstly, the system must have a mechanism to ensure that data integrity is maintained, i.e. there must be a way to make sure that all data held are accurate and can be maintained correctly. Secondly, the data must be protected from accidental damage or leakage. In most environments, these goals can be achieved by ensuring that adequate backups are maintained (including ensuring that backups are restorable) and that data are protected from prying eyes through mechanisms such as strong encryption or physical restrictions placed on the storage devices.

Because data represents something about entities, such as details of bank accounts, it is often seen as the direct target of an attack. Any failings in data security, through flawed procedures, activities, infrastructure, organisation, environment or entities, leaves the system vulnerable.

2.1.8 Application to investigation

By mapping any system onto this seven-element model, it is possible to discover where the weaknesses in the system lie, and hence possible to identify which parts of the system allowed the unwanted activity to occur. It does not, however, always allow us to identify exactly how that activity took place. For this, another model which considers the possible interactions with the system is useful.

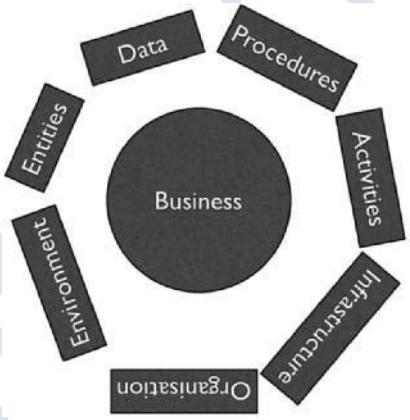

Figure 2.1 A seven-element model of information security

2.2 A developmental model of digital systems:-

In “Silicon pathology” it is proposed that a digital system has a pattern of development similar to the growth and education of a human being. It is also suggested that, in this context, there are four particular routes by which any data or program can arrive on a system. These are summarised in Table 2.1.

This model suggests that any system can be used by two different classes of user – Authorised and Unauthorised. Put simply, authorised users are those who have been granted permission to use the system at the time under consideration, leaving everyone else in the category of unauthorised users.

Table 2.1 Routes by which data/programs can arrive on a system

	Knowing Act	Unknowing Act
Authorised	Authorised user	Authorised user
	deliberately installs	accidentally installs
	(AK)	(AU)
Unauthorised	Unauthorised user	Unauthorised user
	deliberately installs	accidentally installs
	(UK)	(UU)

Although it may seem strange to include consideration of the time of usage in the definition of an authorised user, it can be significant. When we examine a system to discover how it has been used, often the only clue to who might have used it is the presence of confirmation that a particular user identity has been used to log in to the system. If it is not possible, or usual, for the legitimate owner of that user ID to be active at the time in question, further enquiries are necessary to determine if the user was that person, or some other person who has acquired the user ID tokens for some reason. Thus, even though we may have an accurate record of a user ID becoming active on a system, no responsible digital evidence examiner should ever report this as confirmation that that particular person was responsible for the activity. At best they will state something similar to "the activity was associated with the user account".

Misuse of a user ID does not necessarily indicate malicious intent. In spite of attempts to educate users and enforce correct user ID and password usage controls, in line with good information security policies, it is all too common for users to share IDs and passwords in order to make life easier for each other.

This may have something to do with the perceived value of identity tokens such as user IDs and passwords, but a full exploration of this is beyond the scope of this book.

Whatever the reasons for insecurity of ID tokens, it is a fact of life that systems and identities can and will be shared. It is also a fact of life that users can and do make mistakes, leading to the two types of activity (or vectors) defined in the model – Knowing and Unknowing.

2.3 Knowing:-

Vectors classified as knowing define any activity where the user is aware of the consequences of their actions. For example, a user who chooses to install a new word processor, having first spent some time investigating the changes that such an installation will make to their system, would be classified as a knowing act.

This category extends to cover situations where a user chooses to install software or create data using software from a trusted source, relying on the integrity of that source to ensure that there are no adverse or unwanted effects.

However, the knowing category also includes deliberate acts intended to implant illicit software or data on a system, providing the person responsible has made a conscious decision to abuse the system in this way. This implantation may be done for any one of many reasons.

2.3.1 Consumption of illegal material

A user may choose to deliberately download material which breaks the law in any way, ranging from copyright violation to illegality of content (e.g. illegal images of children). Because the act of downloading requires several deliberate acts: finding the download site, selecting the file(s), and selecting the file names under which to save them, the actions can be declared to be knowing. The user in this situation may be either Authorised or Unauthorised and determination of this status can be crucial to successful prosecution.

2.3.2 Implantation of illegal material

Following on from deliberate consumption, we turn to deliberate implantation of illegal/illicit material. Again, the motives for performing this act are varied and can range from simple revenge/extortion through to deliberate abuse of someone else's system to turn it into a distribution node, thus masking the origins of the material and allowing criminals to operate more freely by granting a degree of anonymity. In this instance, it is likely that the user will fall into the Unauthorised category and is likely to be operating covertly at the very least. However, if steps are not taken to thoroughly examine the system for the possibility of Unauthorised Knowing acts of implantation, prosecution may be difficult.

2.3.3 Zombies and Bots

Taking the concept of implanting material on someone else's system for the purpose of distribution leads us to the final example of knowing acts – those of Zombies and Bots. In network security terms a Zombie is usually defined as a computer which has had software installed upon it, allowing a third party to take partial or complete control of that system without the user's permission or intervention. A Bot, meanwhile, is a system which contains software which can give the impression of autonomous operation. Zombie Bots can, and do, exist.

Zombies have been used in the past to set up large networks of com-promised machines in order to launch coordinated attacks against web servers and other systems. Bots, on the other hand, are routinely used to harvest e-mail addresses, send spam,[1] illegally copy web material, acquire credit card numbers, impersonate human beings in chat rooms and many other acts.

The knowing category makes no distinction between the intentions of the person carrying out the installation, mainly because the examination of a storage device which has been subject to data/software deposition probably contains no information whatsoever about the state of mind of the person responsible at the time of the events.

2.4 Unknowing

The case of unknowing installation of software/data deals with the situation where the person apparently responsible could not have reasonably been expected to know or predict that their actions were about to cause damage.This category deals with situations where a user is attempting to perform one action but,

because the tools they are using are flawed or have been compromised, their actions permit some unwanted activity to occur.

2.4.1 Web-site effects

In the author's experience as an expert witness for both prosecution and defence dealing with cases involving images of child abuse, one of the most common questions asked has been "could the images have been placed there as a result of a virus or pop-ups on a web page?". The questioner is really asking whether there is any indication that the images arrived as a result of a knowing or unknowing act, given that the person who was using the machine is already known to be an authorised user.

In the case of web pages it is very easy, by using the HTML "<img=...>" tag amongst other methods, to cause a web browser to download images which cannot be seen by the user. In legitimate sites, this technique is used to get images onto the user's machine so that they are available when the user clicks through to view other pages.

Alternatively, the web site may have been designed to open pop-up or pop-under windows, which appear over or below the desired page, respectively. Typically, this method is used in order to force an advertisement to appear on the user's screen.

Both of these methods allow a third party, the web site manager/designer, to cause material to be downloaded to the user's computer when the user chooses to view particular pages on their site. The user usually has no advance indication that the page they are about to view contains any of the code necessary to cause any of the actions mentioned above.

In these cases, the user is responsible for activity which falls into the AU category, while the person who created the webpages is responsible for UK category activity.

2.4.2 Stowaways

Another possibility is that a user has chosen to install a piece of software, or download a batch of files such as music in MP3 format.

To most users, these actions present no obvious risks, but it is perfectly possible that downloaded files may contain more than they appear to. The act of downloading is a category *AK* activity and this can be demonstrated by the fact that the user must give a location and/or name for the file(s) to be downloaded to. Their interaction with the dialog box which prompts them to provide this information gives them a chance to cancel the download. By giving the information and clicking on "Next", "OK", or "Proceed" they have given evidence of their consent to the download proceeding, strongly suggesting that they have deliberately chosen to accept that material onto their computer. However, where they are downloading something like a

.ZIP archive file, which is a container for many files, they have no way of knowing the exact contents of that file until it has been fully downloaded and uncompressed. If the creator of the compressed file has chosen to include additional files, some or all of which are illegal in nature, then the user has apparently downloaded those files at the same time as the ones they really wanted. The act of downloading is classified as *AK*, the package creator's act is also *AK*, but the user's possession of the illegal files can be defined as *AU*.

2.4.3 File sharing

Consideration of file downloading, as above, leads naturally to the issue of file sharing. Programs such as Limewire, eMule, Kazaa, BitTorrent and Gnutella are designed to allow many users to share complete or partial files to allow them to be transferred more quickly around the community. The software itself has many legitimate purposes, but is also regularly used for illegal distribution of material.

Most file sharing systems are geared up so that the user must choose which files they wish to either download from, or make available to, the rest of the community of users of that particular file-sharing network. Once a file has been downloaded, in full or in part, the default is to make it available to all other users, thus increasing the efficiency of the whole network. Popular files are available from more places. Files obtained from these networks are subject to the same caveats as for any other downloaded file. The person acquiring the file can, almost certainly, have no guarantee that the contents are "as described" until the whole file has arrived.

Some systems, however, operate more like the file sharing found in operating systems. They allow the user to share a folder or part of their storage either as read-only or read/write storage with other users on the network. If the storage area is shared as read/write it is, of course, possible for other users on the network to place files onto it without the owner's knowledge. Although the external users appear to be authorised, whether all of them actually are authorised remains open to debate, as does the issue of which parties are knowing and unknowing **.**

2.4.4 Malware

A final area to consider in this section is that of Malware. Malware is simply defined as software with a malicious purpose. Typically it comprises the family of programs known as Viruses, Worms and Trojan Horses.

A virus, in the digital world, is any program which is capable of replicating itself from one system to another, through some carrier medium, without direct human action. Its propagation is, therefore, UK facilitated by AK because an unknown and unauthorised person deliberately created it, but the action of an authorised person may be required to enable its distribution. In the 1980s, viruses were transferred mainly through the use of floppy discs being swapped between machines. As in the biological world, a virus cannot exist without a host, so viral software spread by infecting other programs, effectively becoming a stowaway within legitimate programs.

In 1988 Robert Morris accidentally created the first Internet Worm. Morris's intention had been to write a program which could calculate the size of the Internet. His aim was to create a program which would hop from machine to machine by exploiting flaws in the security in certain key programs. His design was too successful and contained a serious flaw.

If the worm was only going to explore the Internet, once it had catalogued each machine it should have reported its findings, infected a new system and then died, returning the infected machine to a "clean" state. Morris's worm did not always do this. In approximately one in seven cases, the worm became immortal and continued to send out infections to other machines on the net. Machines which had previously been

infected became re-infected, increasing the chances that they would acquire an immortal version of the worm. The net effect was that within a very few hours, a huge number of machines had been subjected to the attack and the Internet itself was suffering from a "traffic jam" caused by all the worms seeking out more machines to infect.

This attack is clearly of type UK – an unauthorised user has knowingly created a piece of software which places itself onto other systems. However, it could be argued that there is also a UU component to it, as the severity of the attack goes beyond the original intention.

Modern worms follow this type of behaviour, exploiting new security flaws, in order to plant software onto infected machines. Often this software may not cause another machine to be infected, although many do, but may instead plant another program which serves the purposes of a remote criminal.

Finally, we have the group of malware known as Trojan Horses. These are programs which arrive, carried with or by some other program (worms or viruses), and which implant themselves in the system to cause further damage, allow a remote attacker to take control of the machine in order to acquire information from it, or to turn it into a distribution node for e-mail spam, web pages, images or any other files. In these cases, Trojan Horses are now commonly spread through viruses or worms, via infected systems. This mechanism provides the criminal with added security as an investigation needs to trace the origin of the material back to the infected machine and from there back, possibly through a long chain of infected machines, to the originating system. By the time the origin is discovered, it is likely that the criminal will have moved on and have become effectively untraceable.

Since malware has the potential to infect any machine, a responsible examiner should always check for the presence of malware or artefacts indicating prior infection, and report their findings. Most of the major anti-virus software firms maintain databases of malware which de-scribe, in some detail, the effects of each virus, worm or Trojan horse detected.

It must be remembered, though, that anti-virus software tends to operate in the same way as the human immune system. Action to eliminate the invader is only taken after infection has taken place. Anti-malware software can only be created after the malware has been detected and analysed. There will always be a period of time during which successful malware can spread freely.

2.5 Audit and logs:-

Implicit in the discussions above is the assumption that sufficient information about the system exists to allow the models to be applied. Correctly specified systems will have been subject to rigorous specification and testing prior to deployment, and there will be audited records of everything that has officially happened to them in their lifetimes. Unfortunately, in all but the largest organisations, this process tends to be avoided or done badly.

More realistically, the investigator will be dependent on the recollections of the owner/administrator/user about what they think they have done, coupled with any information that the system may have recorded in its own og files. Some systems record copious amounts of data as they are running, and this can be an

invaluable starting point, but others are incapable of recording anything because of the way they are designed or because of limitations on storage space available.

For these reasons, it can be vital for a full audit of a system to be conducted early in an investigation in order to attempt to identify the installed elements and compare with lists of known good or bad components. Time-line analysis, based on the MAC times present is also useful in aiding attempts to identify when system modifications occurred.

2.6 Data content

Data content is, as the name suggests, concerned with the meaning of the data itself. It deals with interpretation of stored data, based on the contents of files.

2.6.1 Data meanings

The starting point for any examination of data content must be the determination of how data can be interpreted. At heart, all digital files consist purely of binary (1s and 0s) data which can be read in many ways (see Table 2.6.1 for examples). As standard, binary digits (*bits*) are read in groups of 8 (8 bits = 1 byte), 16 (word), 32 (double word) or 64 (quad word). The standard used is determined, first of all, by how the processor in the machine accesses memory, but secondly by how software chooses to represent data.

Table 2.6.1 Possible meanings of some byte-length binary strings

Binary String	*One's Complement*	*Two's Complement*	*Unsigned*	*Binary Coded Decimal (BCD)*	*ASCII*	*EBCDIC*
00000001	1	1	1	1	Start of Header	Start of Header
00010001	17	17	17	11	Transmit On	Transmit On
01000000	64	64	64	40	@	Space
01000001	65	65	65	41	A	No-break Space
10000000	−0	−127	128	80	*no meaning*	*undefined*
10000001	−1	−126	129	81	*no meaning*	a

From the processor's perspective, binary data may be program instructions, addresses of data in memory, or just raw data used by programs. This meaning may not become fully clear until the processor starts to work with the data, and the meaning may change at different points in the processing.

Programs usually base their internal data processing on one of the many data encoding standards already defined, but are free to use any representation which their programmers choose to create. For example, an older word processor may use 8-bit data interpreted as ASCII or EBCDIC to represent the text it works with. These character representations can be quite limiting as they contain little allowance, if any,

for the use of different alphabets in different parts of the world. A more modern word processor would use 16-bit Unicode to allow mixed-alphabet documents to be used.

Without any indication of the intended meaning of a binary string, it can be read in more than a dozen ways – each of which is equally likely to be correct.

2.6.2 File extensions, signatures and magic numbers

Because data files can, and are, shared by different programs, mechanisms for passing information about file content have been developed. For the user, extensions to the filename are commonly used, with well-known abbreviations such as ".DOC", ".JPG" and ".PDF" being used to help human beings identify files of similar type.

Table 2.6.2 Sample file signature "magic numbers"

File Type	*Binary Preamble*	*ASCII Signature*
Microsoft Cabinet (.CAB) file		MSCF
for software installation		
Microsoft XBox		XBEH
(.XBE) program		
Portable Network Graphics	01011001	PNG
(.PNG) file		
JFIF JPEG	11111111 01101000	JFIF
(.JPG) file		
MPEG2 layer 3	11111111 11111010	
(.MP3) audio file		

(If a column is blank, there is no data in this format in the header.)

For software, however, the filename has little or no meaning. Instead, most programs use a sequence of bytes at the beginning of each file to determine the nature of the data held in the rest of the file. These file signatures, also known as "magic numbers" are defined for most file types and can be found in reference tables available online (see Table 2.6.2 for a few examples). The signatures are not all the same length and may begin several bytes into the file, using a mixture of raw binary and otherwise encoded data to identify the true purpose of the file.

Many files also have a defined footer or trailer which is used to identify the end of the file and may have special meaning to programs which understand that type of file. Ideally, an examination will use complete files with intact headers and footers, but in extreme circumstances it is sometimes possible to "guesstimate" the original type of file which contained an unattached fragment of data, based on its contents alone. This can be difficult to justify and explain, but does have uses where files have been partially overwritten.

2.6.3 Compression

Many data representations are considered to be inefficient and wasteful of space or network bandwidth. Something like English text, for example

Table 2.6.3 The alphabet in Morse Code

A	.–	B	–...
C	–.–.	D	–..
E	.	F	..–.
G	––.	H	....
I	..	J	.–––
K	–.–	L	.–..
M	––	N	–.
O	–––	P	.––.
Q	––.–	R	.–.
S	...	T	–
U	..–	V	...–
W	.––	X	–..–
Y	–.––	Z	––..

Contains many repetitions and, in ASCII, EBCDIC or Unicode, uses the same number of bits to represent common and uncommon letters alike. In the days when telegraphy was a major communications method, this problem was known and Samuel Morse's code (Table 2.6.3) used shorter patterns of "dots" and "dashes" for common letters to allow operators to transmit messages more quickly. Other innovators developed Morse's principles further by using special codes to represent whole words. In the modern world, we have seen a resurgence of this method in the appearance of "txt-spk" (text-speak, Figure 2.6.1) where abbreviations and altered codings are used to represent common phonemes and words in mobile-phone short messages. All of these methods share a common purpose: they aim to reduce storage space or transmission time by changing the representation of the data without destroying any of the information carried.

Lossless compression

Lossless compression describes a family of mechanisms by which data can be transformed into smaller representations without losing meaning. Detail of some of these mechanisms is beyond the scope of this book, but the principle can be demonstrated with some simple examples.

Consider the picture shown in Figure 2.6.2. This is a 10 × 10 grid in which each pixel (picture element) is either black or white.

Representing this image using the simplest method entails recording every pixel exactly as it appears. For this image we can use one byte value to represent black (giving it the value 11111111) and one byte value to represent white (00000000). Figure 6.3 shows the picture as a 10 × 10 grid of bytes (i.e. it requires a full 100 bytes to represent the whole picture).

Inspection of this binary representation reveals that it is made up of only two bit patterns and that we could make significant savings if we could reduce the number of bits required to represent each pixel. Because there are only two colours, it would make sense to use a single bit to represent either white (0) or black (1), but in order for this mechanism to work for an arbitary number of colours and allow accurate reconstruction of the original image, we need to introduce a colour table into the file. So,

1111111	1111111	11111111	11111111	11111111	1111111	11111111	1111111	1111111	11111111
1111111	1111111			0000000				1111111	1111111
11111111		1111111		0000000	0000000		1111111		1111111
1111111	0000000	0000000	1111111	0000000	0000000	1111111	0000000	0000000	11111111
11111111		0000000	0000000	11111111	11111111	0000000			11111111
1111111	0000000	0000000	0000000	11111111	11111111	0000000	0000000	0000000	11111111
11111111	0000000	0000000	11111111	0000000	0000000	1111111	0000000	0000000	11111111
1111111	0000000	1111111	0000000	0000000	0000000	0000000	11111111	0000000	1111111
1111111	1111111	0000000	0000000	0000000	0000000	0000000	0000000	1111111	11111111
11111111	1111111	11111111	1111111	1111111	11111111	11111111	11111111	11111111	11111111

Figure 2.6.1 "Raw" image in binary form

we can define a file header as follows:

byte 1 number of colours present

byte 2 colour code 1 (bit pattern used to represent colour 1)

byte 3 colour value 1 (original colour in the image)

byte 4 colour code 2

byte 5 colour value 2 . . .

The length of the bit patterns used in the image will be determined by the number of colours present so, in our case, although we are defining the bit patterns in bytes, only the right hand side (least significant) bit will actually be used as we only required a single bit for two colours. Applying this, our 10 × 10 grid becomes that shown in Figure 6.4 and occupies just 100 bits which is just 12 $\frac{1}{2}$ bytes.

Even with the addition of the file header, which contains the details of the original colours in the image and the bit patterns used to represent them, the compressed file is still only 17 $\frac{1}{2}$ bytes in length, considerably smaller than the original, but without losing any of the original information. of course, the example given above is a little artificial and does not work for all images, but it adheres to the main principle of lossless compression: no information has been destroyed. Its representation has just been transformed into something more efficient.

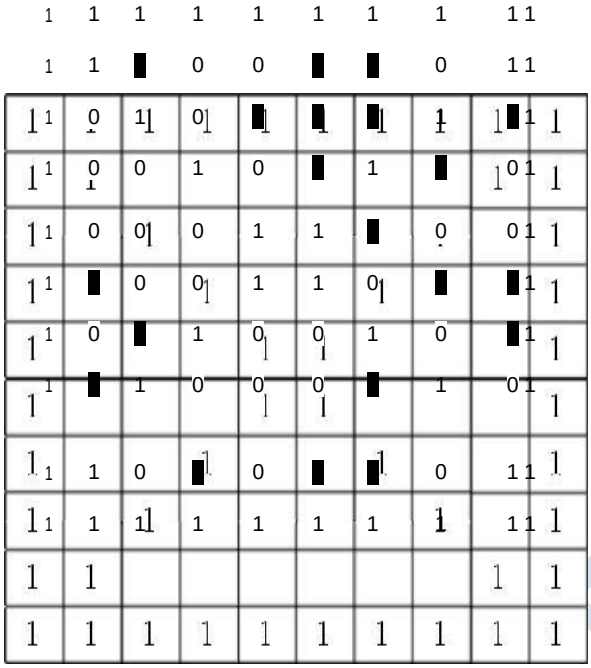

Figure 2.6.2 Compressed image using one bit per pixel

Compression of this type is particularly appropriate where damage to the original information cannot be tolerated, and can be applied to files such as word processor documents, databases, spreadsheets and financial records.Depending on the nature of the data to be compressed, a range of more advanced lossless compression techniques may be used and the reader is advised to look at Lempel-Ziv (LZ) and Huffman as good starting points if more detail is required.

Lossy compression

Lossy compression takes compression a stage further, accepting that not all information is always required. In photographs, video and audio in particular there is often much detail which goes unnoticed by the human eye, or which is imperceptible to the ear. If these "unnecessary" components can be identified and discarded, an acceptable impression of the original is retained, but in a much smaller file which can be further compressed using lossless techniques.

When the file's contents need to be reproduced, the software will inter-polate, filling in the missing detail based on the information remaining in the file.For most purposes, this results in an acceptable representation of the original in spite of the loss of information. However, if the compression

Figure 2.6.3 Original uncompressed image: file size 320 kilobytes

algorithm is forced to attempt too high a level of compression, too much in-formation will be deleted and the decompression software will be unable to produce a good replica of the original. Figures 2.6.1, 2.6.2 and 2.6.3 show JPEG images with varying levels of compression. Figure 2.6.3, in particular, shows how badly images can be degraded by over-compression using lossy methods.

As the images show, the loss of detail between Figures 2.6.1 and 2.6.2 does not significantly impact our perception of the overall image, although careful inspection shows some loss of detail in the buildings and water. Figure 2.6.3, on the other hand, is unrecognisable, showing unacceptable degradation.

Historically, this has been a source of debate about the admissibility and use of compressed digital photographs for forensic purposes, but it is now accepted that, except in situations where any detail loss is unacceptable, they provide an adequate representation of the object being photographed and can be used in the same way as conventional "wet film" photographs.

Figure 2.6.3.1 Image compressed using JFIF JPEG set at 50 per cent quality: file size 24 kilobytes

2.6.4 Composite files

In the early days of computing, application programs tended to work with only one type of data and required only simple data files which contained all the information appropriate for the task in hand. These monolitihic files provided only simple storage capabilities, but were adequate for their intended purpose and many applications are still capable of working with those simple files.

However, as we have moved towards more mixed or multimedia applications, files have become more complex to the point where modern data files used by word processors, presentation programs, spreadsheets etc. may have to hold a mixture of numerical, textual, graphical and audio data all at the same time. These "composite" files typically exist in two forms: Segmented and Compressed Folder.

A segmented file appears, superficially, to be a monolithic file, but is usually a sequence of data "chunks", each one containing a different part of the data required for the document in question (see Figure 6.8 for an example). This allows a single file to contain different media types in a manner which allows all the individual elements of the document to be kept together for storage and transfer.

More recently, though, programmers have started to explore the possibilites of XML , the extensible Markup Language, as a way of producing

Figure 2.6.4 Image compressed using JFIF JPEG set at 10 per cent quality: occupies 4 kilobytes

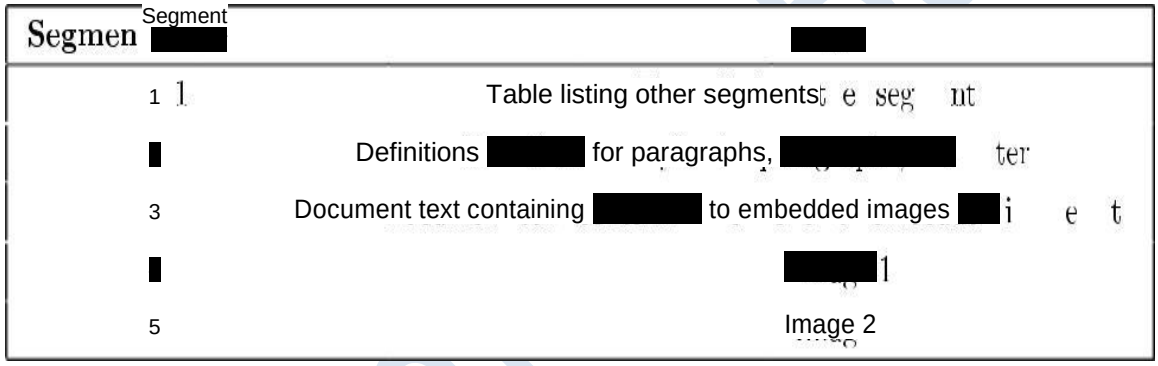

Figure 2.6.5 An idealised sample segmented file structure for a word-processed document

catalogues of the elements which need to be combined to produce a document. They have created an expanding family of composite files, such as the Open Document and DOCX standards, which mainly consist of a compressed directory or folder holding all of the document elements. Within the directory, the document elements are held as individual files, along with a XML "manifest" file detailing the nature and purpose of each of the files in the directory.

The net effect is similar to that achieved by the segmented file, but modification to the composite file can be made more easily, as each component can be edited individually. Moreover, because of the use of XML, there is potential for new features to be added in a way which does not prevent older versions of the programs from handling newer files correctly.

2.6.5 Encryption

The use and even the very existence of encryption technology has provoked much debate and controversy over the years. There are those who believe that it should have no place in a civilised society and that no-one should be permitted to use it, while others accept that it offers an essential means of protecting privacy and freedom of speech. Whatever your point of view may be, the reality is that strong encryption systems exist and are freely available to anyone who cares to buy or download them. As a result, some criminals choose to use encryption in an attempt to disguise their activities.

Under the terms of the Regulation of Investigatory Powers Act usually shortened to RIPA, failure to disclose encryption keys is an offence punishable by up to two years in prison. Obviously, if the offence under investigation carries a longer sentence it may be preferable, from a criminal's perspective, to be found guilty of the RIPA offence whilst preventing any evidence of a more serious offence being uncovered.

The goal of any encryption system is to scramble the original data or message in such a way that no unauthorised person can decrypt it to recover the original content. This is usually done by ensuring that only authorised

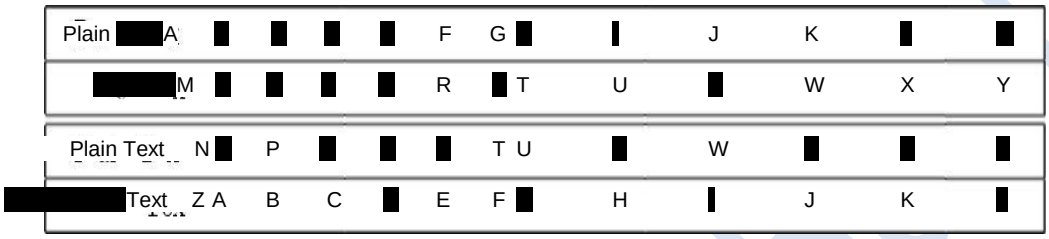

Figure 2.6.6 A simple shift-substitution or "Caesar" cipher table

Recipients have access to the encryption "key". This key takes many forms: it may consist of knowledge of the actual encryption method use; it may be a secret word, phrase or number needed to decode the data; or it may even be some physical feature of the recipient (such as a fingerprint) which can be used to generate a unique key for that person.

Over the years, encryption has developed a long way from the simple substitution cipher (see Figures 2.6.6 and 2.6.7), often credited to Julius Caesar, where each letter of the alphabet is replaced by a single other letter, thus obscuring the message's contents to the casual observer.

Simple ciphers, such as the Caesar cipher, are susceptible to easy decipher-ing using statistical cryptanalysis methods based on knowledge of the language thought to have been used to construct the original message. In our example (Figure 2.6.7), the letter "E" appears most often, as it does in written English, and we could deduce that it has been substituted by "Q" in the ciphertext. This leads us to the key that A becomes M, B becomes N and so on.

If the key is not quite as obvious, we could perform further analysis based on knowledge of letter frequencies in written English and deduce each letter used in the message individually. Couple this with our knowledge of how words are spelt, and a few minutes of effort would break this cipher easily.

More complex ciphers aim to break the patterns present in the data by using more complex keys and manipulating data at lower levels. Typically, a modern encryption system will use a key of at least 1024 to 2048 bits as the basis for its encryption. By combining this key with the original binary data, the resulting encrypted data can appear to be nothing more than random gibberish. However, the very randomness of the resulting file can at least assist us in detecting the presence of encryption through an "entropy" test.

Figure 2.6.7 Application of the cipher table in Figure 2.6.6

Files with high entropy appear to be highly random and, as noted above, this is often a good indicator that encryption has been applied. If we can determine the software used and obtain the key, possibly using the methods described in Chapter 4, then the encrypted data can be recovered.

If we are unlucky, though, it may be necessary to attempt a "brute force" decryption attack where many different algorithms and keys are tried until the data are recovered. The complexity of modern encryption means that such attacks may take several months or even years to complete successfully, if at all.

2.6.6 Steganography

An alternative, or adjunct, to encryption is the use of steganography, or data hiding, to disguise the presence of material. This concept, too, predates the computer age, and its existence can be traced back for several centuries. Allegedly, one of the early examples of steganography involved shaving a man's head and then tattooing a map on it. Once the hair had grown back, the existence of the map was hidden from sight and only those who knew it was there would seek out the carrier.

In more recent times, secret agents have reduced photographs to the size of a full stop on this page and stuck them on pages of text. Again, only someone who knows where to look is likely to find information hidden in this way.

The content of a file may, itself, contain hidden meaning. For example, a terrorist leader who releases a video clip to television news channels may have arranged various "props" in the background and the presence and positioning of those props may communicate a message. The ordering of items in a shopping list, choice of music or desktop wallpaper may also be used in this way.

However, for our purposes, we shall concentrate on digital steganography. Modern steganographic methods rely on properties of files to allow data to be combined to produce an innocuous carried file which contains the secret information. As outlined in the earlier parts of this chapter, digital files have specific structures and may, as a result, have areas which are under-used or amenable to modification without significantly affecting their use.

Most current research on steganography is concerned with methods for embedding, and detection of embedded content, in files produced through lossy compression. Because these files are already distorted representations of the original, it is possible to make minor modifications to them without affecting the user's perception of the reconstructed image, movie or sound. As an example, consider a single pixel in a photograph. Effectively, within the file, information is stored about the colour present in each pixel. This may be as an absolute colour value, typically using eight bits for each of Red, Green and Blue (giving an RGB triple) where 00000000 00000000 00000000 is pure black, 11111111 00000000 00000000 is Red and so on, to 11111111 11111111 11111111 giving pure white. Alternatively, it may be stored as a value which represents the colour of that pixel relative to those around it (e.g. one unit "bluer", three units "red-greener" etc.).

No matter what the representation used in the file is, the ultimate goal is to allow it to store something like the original information in a way which the user perceives as "good enough". Making changes in the low-

order bits (i.e. those which have least impact on the resulting value) does change the reconstructed image, but in a manner which is almost imperceptible (00000000 00000000 00100010 is not a very different shade of Blue to 00000000 00000001 00100010). Thus a single byte of hidden content can be distributed across several pixels, making a very small change in each.

So, the human being will not perceive the changes, particularly if they do not have access to the original information, but appropriate software which knows the embedding mechanism can extract the hidden data on demand.

Detection of this type of embedding, currently, tends to rely on either knowing the signatures of common embedding programs (i.e. common patterns which always seem to appear where these tools have been used), or statistical analyses of files which measure deviation from the normal profile of known "clean" files.

Other possibilites exist in segmented and composite files. It is entirely possible that additional data can be added to these files, but since it is not used by the reading program, it will be invisible to the regular user. In this case, detection can be performed by careful inspection of the segment table or the manifest and identifying those parts of the file which are clearly unused in the carrier document. Even when steganographic content can be detected, though, it is often found to be encrypted . . .

2.7 Data context

While there is little doubt that, in terms of initial evidence, file/data content is crucial, all that it really provides is evidence of the existence of those data on the device being interrogated. In order to determine something more about how the data arrived and what they mean, we need to consider the context in which they exist.

Every file on a system should have timestamps associated with it, detailing when the file was created, last modified and last accessed. Consideration of these timestamps allows us to build up a picture of the sequence of events during a period of usage. The closer that period of usage is to the time of seizure, the more accurate or complete the sequence of events will be. Also of significance is the exact name and location of the file within the filesystem.

By default, many programs automatically create data files in special areas of the filesystems which are, effectively, reserved for their use. Usually, these "reserved" areas will be located in common areas which the user would not normally be expected to explore.

Files found in these areas, therefore, tend to indicate that a particular program was in use at the times indicated by the files, but may indicate nothing more than an automatic action of the software.

Files found in other areas, however, such as the "My Documents" folder tree in Windows, tend to have stronger evidential value. For a file to appear in these areas, or for it to have a name which is obviously meaningful to a human being, tends to indicate that the user has made a conscious decision to create and manipulate that file, with that name in that location. This meta-data is not necessarily preserved when the file is deleted, but it will always exist for live files.

Thus it is possible to have lengthy discussions and debate about not only the nature of data content, based on how different programs may interpret the data, but also on the meaning of the meta data and the significance of a file found in one location as opposed to a copy of it found in another.

2.8 The ISO/OSI model

The ISO/OSI model divides the functions of a network into seven layers based on the services each layer provides and requires (see Figure 2.8.1).

Each of the seven layers is concerned with a single key function and provides a service to the layer above, while using the services provided by the layer below. The great benefit of this type of model is that the network can adapt to embrace new technologies by fitting them into the appropriate layer(s). If done correctly, there is little or no impact on the layers above or below and software can exploit the new functions with no modifications at all.

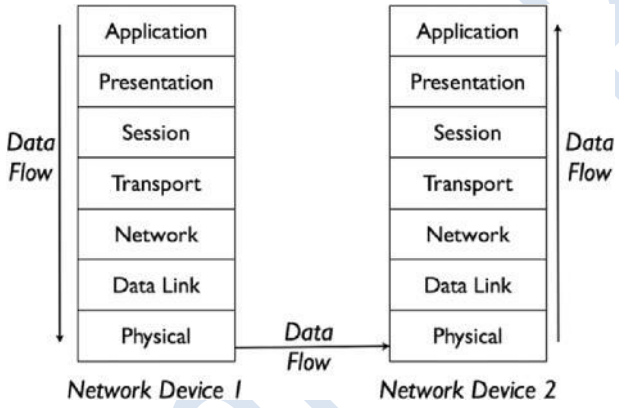

Figure 2.8.1 The ISO/OSI seven-layer network model

2.8.1 Application layer

The application layer is the highest level in the network stack and represents the languages used by programs to communicate with each other. The protocols in this layer tend to represent discrete actions required by the software to support specific operations and, on the whole, are designed to support client-server model activities where one program (the client) requests a service of the other (the server). Some systems operate in a pure client-server model, where it is obvious which component is the server as it completely controls access to its resources. Others, however, operate more like peer-to-peer systems, where all programs provide both client and server functionality.

A client-server system is sometimes described as "fast-food computing". In order for a customer (client) to get access to the food, they must issue a request to the shop assistant (server) who will fetch it for them, or make them wait until it is available. If the server is busy, the clients have to queue and await their turn.

2.8.2 Presentation layer

The presentation layer defines the network equivalent of the alphabet to be used for communications. That is to say, it defines the individual language elements which are used to represent the application layer protocol. This layer can be responsible for translating between different character sets (e.g. ASCII to

EBCDIC and back again) and may also be responsible for implementing the encryption used to provide secure communications across an untrusted network.

2.8.3 Session layer

Moving down the stack we next encounter the session layer which is responsible for maintaining a dialogue between two pieces of software. This layer is the lowest one at which we can distinctly identify individual conversations. The session layer separates traffic passed to it by the transport layer and ensures that the various fragments of "conversation" are sent to the correct recipient. As an analogy, it can be thought of as the person who delivers letters to the correct offices from a central mailroom that receives all post for the building. It has a particularly important role at the beginning and end of any dialogue as this layer is responsible for ensuring that the communications session is established and closed correctly.

2.8.4 Transport layer

While the session layer is concerned with program-to-program communications, the transport layer deals with end-node to end-node (or source machine to destination machine) communications, and is responsible for combining multiple sessions into a format suitable for transmission between machines, making best use of the available communications bandwidth.

To do this, the transport layer typically breaks up session layer data into chunks, known as "packets" which can travel independently of each other. At the receiving end, the corresponding transport layer gathers packets together and reassembles the data contained in them into something which can be passed to the session layer.

It is a little like a courier travelling from one building to another. All the letters produced by the session layer are handed to the courier and delivered to the correct building. At that building, the transport layer (courier) hands the letters over to the session layer (mailroom) for delivery to the correct office.

This layer can usually detect simple communications failure and deal with non-receipt of data by requesting re-transmission from the corresponding transport layer at the other end of the communications channel.

2.8.5 Network layer

The network layer is the route planner of the stack. Its concern lies with identifying the most efficient way of getting the transport layer's data from point of origin to destination. This may involve pre-planning a route and attaching information about it to the data before it is sent, or it may involve individual intermediate nodes making their own decisions about the best route on a case-by-case basis as each chunk of data is processed.

Although the transport layer views the connection between two machines as a continuous channel full of packets, the network layer may implement it as a series of "hops", passing data from one machine to another until the final destination is reached. The exact route taken by each packet may change based on network conditions, or it may be the same for all, depending on how the network layer has been told to behave. The network layer also attempts to monitor the state of those parts of the network that it uses, and

adjust its routing tables appropriately to suit prevailing conditions, unless instructed to use static routes for certain communications.

2.8.6 Data link layer

Sitting between the network layer and the physical layer is the data link layer. It is responsible for implementing the network hops used by the network layer, and deals with adjacent node-to-node communications. As such, it is aware of the real structure of the network as created by the physical implementation, and often uses a different type of address to identify a particular piece of hardware. Since this address is the one used on the communications medium and which, effectively, controls access to the data link layer, it is usually known as the Medium Access Control address. For most devices, this address is set at the point of manufacture and is not normally changed without very good reason.

The data link layer also repackages data from the network layer into chunks (or frames) which are more suitable for transmission by the physical layer.

2.8.7 Physical layer

At this lowest level of the network stack, the data link layer frames are converted into something which can be transported through the medium which carries the network. This can be electrical impulses for wired net-works, light for optical networks or radio frequency signals for wireless networks.

2.9 The Internet Protocol suite

The development of the Internet Protocol (IP) suite, in the late 1960s and 1970s is possibly one of the most significant events in human history. Until this point, all networks tended to be tied into particular manufacturers and hardware, but the IP suite, with its layered model, created a de-facto standard which everyone could, and did, adopt. A large part of the popularity of this technology also comes from the fact that all significant developments have been subjected to public discussion, scrutiny and publication through the RFC (Request for Comments) mechanism adopted in the early days of its development. Through this system, anyone was free to propose a new standard for internetworking and subject it to inspection by the rest of the community. This, in a manner similar to the development of modern open-source software, led to the creation of robust standards which could be implemented by all.

Nowadays, the RFC mechanism is still used, although in a more formalised manner, and RFCs are archived online in various locations such as http://www.rfc.net/ and http://www.rfcs.org/.

Where the ISO/OSI model proposes seven layers, the IP suite contains only five (see Figure 7.2), merging or distributing some of the ISO/OSI functions within those layers.

2.9.1 Application layer

In the IP suite, the application layer combines the functions of the application and presentation layers in the ISO/OSI model. At this level we find protocols such as HTTP (the Hypertext Transfer Protocol used in most common WWW applications), SMTP (Simple Mail Transfer Protocol), POP-3 (Post Office Protocol 3), FTP (File Transfer Protocol) and all the

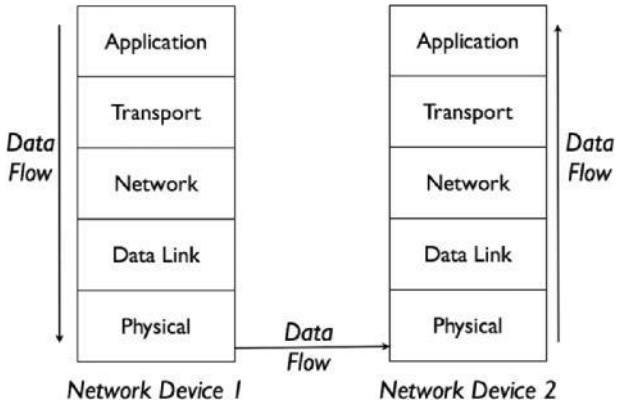

Figure 2.9.1 The IP suite five-layer model

Others regularly used by individual pieces of software. Some of these are discussed in more detail later in this chapter.

2.9.2 Transport layer

IP's Transport layer shares some of the functions of the session and transport layers in the seven-layer model. Two main protocols are used at this level: the Universal Datagram Protocol (UDP) and the Transmission Control Protocol (TCP).

UDP, sometimes also known as the Unreliable Data Protocol, places emphasis on rapid delivery of data from one point to another, but gives no guarantee that all data will arrive as intended. It is particularly useful for applications where large quantities of data need to be moved quickly, but where quality of service is not a particular issue. Typically, streaming audio or video services, such as Internet radio or TV, will use UDP as the occasional hiccup in sound or pictures can be tolerated.

TCP, on the other hand, includes mechanisms for guaranteeing that lost packets will be re-transmitted once their loss has been detected. Therefore, although it can provide no particular guarantee about speed of delivery, it is favored for applications where there needs to be a guarantee that all data will be received completely and correctly.

In both these protocols, the source and target applications at each end are identified by unique port numbers. Each protocol has 65536 ports available and they are probably easiest to understand if thought of as extension numbers in a telephone system. Calling the central switchboard ensures that the right company has been reached, asking the operator for a particular extension makes sure that the session is then handled by the right person or department.

2.9.3 Internet layer

The internet layer is broadly equivalent to the network layer of the ISO/OSI model, providing network management and routing services. It is in this layer that we encounter the legendary "IP address" which every machine on the Internet needs in order to be able to send and/or receive data.

At the time of writing, the vast majority of IP addresses in use are from the original IPv4 system, although most operating systems now support longer IPv6 addresses, which were introduced to deal with the

impending shortage of unique addresses and allow for greater expansion of the Internet. At heart, both addressing systems use similar mechanisms to assist traffic routing, so we will concentrate on IPv4 as the simpler of the two to describe.

A standard IPv4 address is actually a single 32 bit number but, for ease of use, is usually written as a sequence of 4 bytes, known as a "dotted quad", as in 127.0.0.1. Historically, three main blocks of IP addresses were available for public use, allocated as:

Class A networks beginning with 0.x.y.z to 127.x.y.z (i.e. 0.0.0.0 to 127. 255.255.255). Effectively this class consisted of 128 large networks, with the network identified by the first byte, and the machine within that network by the remaining three bytes.

Class B networks beginning 128.0.y.z to 191.255.y.z (i.e. 128.0.0.0 to 191.255.255.255). These networks were smaller than Class A and identified by the first two bytes, with the remaining two being used to identify the node within the network.

Class C networks beginning 192.0.0.z to 223.255.255.z (i.e. 192.0.0.0 to 223.255.255.255). In these networks, only 255 nodes are available (from the last byte) with the leading three bytes being used to identify the network.

Within these networks, three special networks were reserved for experimental or internal use and their addresses should never appear on the public Internet. These are the networks beginning 10.x.y.z, 172.16.y.z and 192.168.0.z to 192.168.255.z. Several other blocks of addresses are reserved for other purposes, including the 127.x.y.z network which always represents a virtual (non-physical) network that exists inside each and every Internet-capable machine.

The use of one part of the address as a network identifier, with the remainder as the node identifier, simplifies the job of delivering data from one place to another by splitting the delivery task into two stages. The first stage uses the network prefix to deliver data to the correct network; once this has been accomplished, the receiving network can use the rest of the address to ensure that data is sent to the right node.

Address allocation under the "classful" scheme shown above is recognised to be wasteful, so most modern networks now use a modified system called CIDR (Classless InterDomain Routing) addressing to allow any number of bits to be used for network identification with the remainder of the 32 used for node identification. Under this system, addresses are usually written as a dotted quad followed by the number of bits used for the net-work ID (e.g. 172.16.8.33/12 means that the first 12 bits are the network address and the remaining 20 are the node address).

2.9.4 Data link layer

This is the equivalent of the seven-layer version's data link layer, and is still responsible for node-to-node communications as directed by the Internet layer. The exact detail of this depends on the underlying physical layer but for most Internet systems there will be either a MAC address, associated with a network card, or a telephone number associated with the physical communications channel. Somewhere in the network there will be a record of the relationship between IP address and data link layer identity. However, it should be noted that, since IP addresses can be allocated on demand and thus systems may have different

IP addresses at different times, it is vital to know the exact time that an IP address was in use if it is to be traced accurately.

2.9.5 Physical layer

Finally, the physical layer continues to convert data link layer frames into appropriate signals for the transmission medium.

2.10 DNS

Although, at machine level, the Internet relies on the use of IP addresses, these are less than intuitive and not particularly useful for most humans.To solve this problem, a system to allow humans to refer to machines by name was introduced. This is is known as the Domain Name System/Service or DNS. In DNS, machines are grouped into a hierarchy of names, allowing related machines to be identified easily while giving enough flexibility for each machine to have several names and/or IP addresses associated with it.

At the top of the DNS hierarchy lie 13 "root" servers responsible for the notional "." highest level domain. These servers contain the details of the servers for the approved top-level domains (TLDs) such as "com", "gov", "mil", "uk", "tv" etc.Each TLD has its own group of servers which contain pointers to the servers for the next layer down (e.g. ".ac.uk", ".co.uk", "microsoft.com") and so on until the servers at the bottom of the tree contain only details of the relationship between complete "Fully Qualified Domain Names" (FQDNs) for particular systems on the Internet, and their IP addresses.

Each network should have at least one DNS server to handle queries from its own machines, answer queries about its own domains or both.When a machine needs to translate between a FQDN and the corresponding IP address, it queries its own DNS server (the local server) which either responds with the address (if it already knows it) or queries a more authoritative server for the domain. If the domain cannot be found, queries will be sent to one of the root servers which will direct the local server to the appropriate TLD server which directs the query down to a lower level, and so on until the master server for the domain has been reached. Once a server has received a response to a domain query, it stores a copy of the response for a set period of time (determined by the domain master) for future use. See Figure 2.10 for an example of this process.

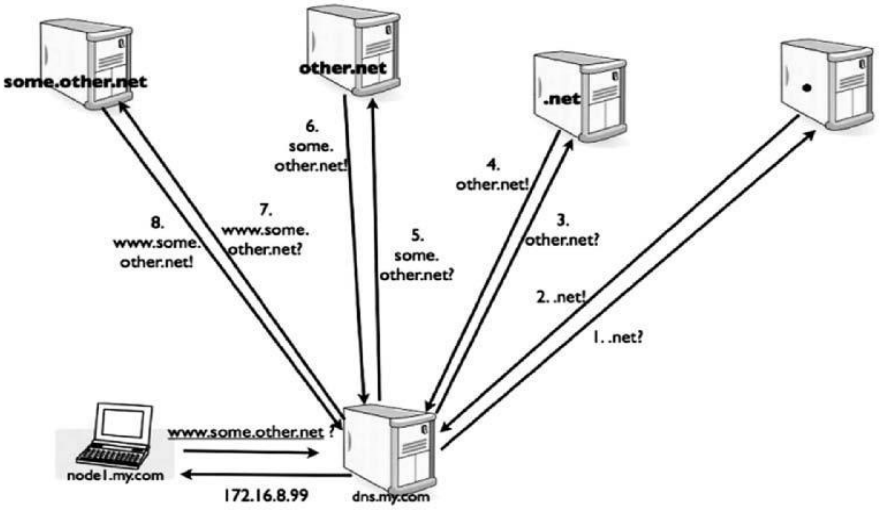

Figure 2.10 A DNS query in progress

In order to have the right to put entries into DNS, someone must claim ownership of the domain. This usually involves buying, or more accurately leasing, a domain name from a registrar recognised by DNS's governing body – the Internet Corporation for Assigned Names and Numbers.[2] Once the rights to use the name have been established, a record of "ownership" of the domain name is placed into one of the international "WHOIS" databases which list all domain names and their owners. Only then can DNS entries be created in the appropriate DNS servers to allow the domain name to be associated with the physical server(s) which host the services it offers. For any enquiry, the WHOIS and DNS databases can provide significant information about the claimed ownership of any domain name and the associated IP addresses.

2.11 Internet applications

2.11.1 WWW

The World Wide Web was created at CERN in 1989 by a British physicist, Tim Berners-Lee, attempting to solve the problem of keeping documents up to date within a large distributed organisation. His work drew on previous experiments with HyperText systems and distributed in-formation systems such as Gopher , but extended them to create a rich "information space" in which different media types and services could be accessed through a common user interface. The browser concept was quickly adopted and developed further by companies such as Netscape to produce a platform which was easy to use and thus appealed to a broad range of users with a wide range of knowledge and experience.

Strictly speaking, therefore, the World Wide Web is really a collection of services sitting on top of the wider Internet and accessible through a unified interface program. To most, though, the widespread use of the web has led to it becoming synonymous with the Internet as a whole, although there are many services available on the Internet which are not accessible through web browsers.

For the purposes of this discussion, we will focus on the two most popular components of the WWW: HTML and HTTP.

HTML is the HyperText Markup Language used to produce the web pages with which we have all become familiar. It is a way of describing the properties of elements on a page, allowing files to be embedded and links to other content to be defined using a simple scheme of "tags" which are written using a simple language designed to be as accessible as possible to non-programmers. A sample of HTML can be seen in Figure 2.11.1

All web pages consist of at least a page of HTML which may contain references to other resources such as hyperlinks or embedded objects. Hyperlinks are defined, within the HTML source, by <a href=...>

```
<html>
<head>
<title>A simple html example</title> </head>
<body>
<p>This is some text in the HTML example and the next bit will embed an
image</p>
<img src=''images/myimage.jpg''/> <p>Now this line contains a
hyper<index>link</index>link to
<a href=''http://www.google.com''>google.com</a> </p>
</body>
```

```
</html>
```

Figure 2.11.1 A sample of HTML used to compose a web page

tags, whereas embedded objects may be defined by tags such as <img src=...> for images and <object> for other objects such as embed-ded programs and movies.

Many of these tags contain references to data stored on the same server as the web page being interpreted, but some references will also be made to content held elsewhere. In these cases, the resource will be identified by a URL[3] which has the general form:

<scheme>://<user>:<password>@<server>:
<port>/<resource-path>?<data>

where <scheme> is the protocol to be used to access the resource (e.g. http), <user> and <password> are the logon details to be sent with the request, <server> is the FQDN or IP address of the server to be contacted, <port> is the TCP port number on which the web server is listening (the default value is 80), <resource-path> is the full name of the resource as it will be understood by the server and <data> is any data to be sent to the server as part of the request. With the excep-tion of <scheme>, <server> and <resource-path>, all of these elements are optional and may be omitted from the URL.

At its simplest, then, a URL needs to be nothing more than

<scheme>://<server>/<resource>

as in

http://www.google.co.uk/

In this example, the <resource> is the shortest one possible – a single /, denoting the default resource.

The default communications protocol for WWW transactions is the HyperText Transfer Protocol, HTTP. This is a very simple language defined in terms of basic operations required to allow information to be exchanged between two machines. By definition, it uses plain text for data presentation, but a variation known as HTTPS (for HTTP-Secure) provides encryption of data between client and server. Where the default port for HTTP is TCP port 80, HTTPS is usually found on TCP port 443 on the server.

Both versions have the same main functions which are defined by a few major commands:

GET: the main HTTP command retrieves content from a specified resource.

HEAD: retrieves information about a resource, without fetching any of the content. The HEAD data is also returned as part of the response to GET.

POST: used to send data back to the server for processing.

PUT: used to upload files/new content to the server.

Web browsers

When a web client (browser) connects to a server to retrieve a web page, it typically starts by issuing a GET command to retrieve the main HTML page. Once the HTML has been received, the browser will then

issue a series of GETs and/or POSTs to retrieve embedded resources identified by the HTML. Thus a single web page, made up of multiple files, will require several requests from client to server.

Because Internet connections used to be very slow, early browsers were programmed to store copies of previously used resources locally in an area known as the *browser cache*. Not all material is cached, particularly if it has been accessed over a secure (HTTPS) connection, but even today most browsers still maintain a cache of previously accessed resources, which can prove invaluable to investigators looking for evidence of web activity. Files in the cache are no different from ordinary files, except that they have been created as a result of automatic download.

The exact location and name of this cache depends on the browser in use and the operating system. Users can also change the location of the cache by modifying browser options. Default cache folders tend to be well hidden, deep down in the user's filespace in areas that are notionally reserved for system and program use.

Most browsers also tend to keep a "history" of recently visited resources, so that popular or recently visited pages can be found from a "drop-down" list rather than forcing the user to remember a full URL. This history is usually stored in a file, somewhere near the cache, and includes date and time of last access. Again, this can be a useful source of information for an investigator.

The final piece of information which can routinely be recovered from web-browser storage is the collection of cookies built up over time. Cookies are small chunks of data stored by the browser, at the request of the web server. They are often used to prove that a user has logged into a web site, store details of items in a shopping basket, remember preferences set by the user, or even just to allow the web server to recognise returning users for statistical purposes. Without them, most of today's web applications would not work.

If the browser accepts and stores a cookie, it then presents the cookie data back to the server as part of subsequent requests. The server can then check the returned value against its own list of known cookies and perform appropriate processing.

Figure 2.11.2 shows the contents of typical cookies. As with caches and history files, different browsers may use different storage mechanisms, but Will store cookies unless the user elects to override this behaviour. Within the cookie, we find the following fields:

	CH 7 INTERNET ACTIVITY				
Created	Domain	Expires	Name	Path	Value
231613802 99180999	scholar.google.co.uk	2038-01-17 19:14:07	GSP	/	ID=5406F79
229273537 89834601	eu.wiley.com	2010-04-07 15:05:37	utma	/	14086598...
229273537 89915299	eu.wiley.com	2008-10-07 03:05:37	utmz	/	14086598...

Figure 2.11.2 Contents of typical cookies from the author's web browser

Created: the time and date of creation – encoded as the number of seconds since a reference date (usually midnight on Ist January, 1970).

Domain: the domain to which the cookie should be returned. This may be a FQDN or an upper-level domain.Expires: the last valid date and time for the cookie. Some cookies may have infinite life, indicated by either a 0 or a large value here.

Name: the name of the cookie as it is known to the server.

Path: the starting point in the resource path for the server. When the browser accesses any resources at or below this path, the cookie will be sent to the server.

Value: the data to be stored in the cookie. This may be numeric or alphabetic, and the meaning may not be obvious to anything except the receiving server.

Because cookies are stored in the browser's storage area, often in plain text or easily decodable files, they can be manipulated and deleted quite easily. In the author's experience, however, it is rare for someone to go to the trouble of manipulating cookies unless they wish to attempt to gain access to a restricted web resource which is protected by a user authentication system and uses cookies as proof that a correct login has occurred. In this situation, the criminal may copy a cookie from a legitimate user and then install it into his/her browser in order to impersonate that legitimate user.

2.11.2 E-mail

Electronic mail is still one of the most popular applications on the Internet and, for most people, can be split into two distinct operations: sending and collecting.

Sending e-mail

At the heart of Internet e-mail lies the SMTP (Simple Mail Transfer Protocol), along with standard definitions of the format of mail messages. These two standards laid the foundations for modern e-mail and, although there have been several refinements documented in later RFCs, they still enshrine the main principles of operation of e-mail sending on the Internet today.

SMTP defines a mechanism which allows mail software to either deliver messages directly to the receiving system, especially if it is local, or to pass undeliverable messages on to "smart" hosts which act as intermediate relays. These smart hosts may also pass mail along to another relay, creating a chain of mail hosts, until the message eventually arrives at a mail server that can deliver the message to the target system. In DNS, most domains have relays defined through the use of Mail eXchange (MX) records, and these machines are usually the only ones which can be reached via port 25, the SMTP port.

As the mail passes through each relay, the relay adds an extra line of information to the top of the message, giving details of where it was received from, the time of receipt and where it seems to be going. In this way, each e-mail transmitted over SMTP picks up a set of "headers" which contain complete details of its journey across the Internet (see Figure 2.11.3).

By default, SMTP relays will hold e-mail and attempt to deliver it for a period of up to seven days before deciding that the receiving network is unreachable and returning it to the point of origin with an error message. Of course, if the receiving network does not exist (i.e. has no valid DNS entry or IP address), the server will reject it immediately.

Amongst the weaknesses in the definitions of SMTP and mail message format is the requirement for mail clients (sending software) to generate

```
Delivered-To: xxxxxxxxxxxx@gmail.com
Received: by 10.142.224.1 with SMTP id w1cs78057wfg; Wed, 5 Mar 2008 16:56:32 -
      0800 (PST)
Received: by 10.82.182.1 with SMTP id e1mr6469630buf.21.1204764990299; Wed, 05 Mar 2008 16:56:30
      -0800 (PST)
Return-Path: Freemanlocomotionhurley@dispatch.com
Received: from xxxxxxx.xxxxxxxxx.co.uk (xxxxxxx.xxxxxxxxx.co.uk [212.67.202.165]) by
      mx.google.com with ESMTP id 6si1305423nfh.30.2008.03.05.16.56.29; Wed, 05 Mar 2008
      16:56:30 -0800 (PST)
Received-SPF: neutral (google.com: 212.67.202.165 is neither permitted nor denied by best guess
           record for domain of
           xxxxxx@xxxx.com) client-ip=212.67.202.165;
Authentication-Results: mx.google.com; spf=neutral (google.com: 212.67.202.165 is neither permitted nor
                  denied by best guess record for domain of xxxxxx@xxxx.com)
                  smtp.mail=xxxxxx@xxxx.com Received: from [76.76.168.165]
                  helo=hp22952133567) by xxxxxxx.xxxxxxxxx.co.uk with smtp (Exim
4.54) id 1JX4PM-0006QG-St
for xxxxxxx@xxxxxx.net; Thu, 06 Mar 2008 00:56:29 +0000 Message-ID:
149b501c87f462b28c6b00a01a8c0@HP22952133567 From: "Myles Cervantes" xxxxxx@xxxx.com
To: <xxxxxxx@xxxxxx.net>
Subject: your exclusive watches rolex Date: Wed, 5 Mar 2008 20:47:56
+0800 MIME-Version: 1.0
Content-Type: text/plain; format=flowed; charset="iso-8859-1";
reply-type=original Content-Transfer-Encoding: 7bit X-Priority: 3
X-MSMail-Priority: Normal
X-Mailer: Microsoft Outlook Express 6.00.2800.1106 X-MimeOLE: Produced By Microsoft
MimeOLE V6.00.2800.1106
Discover Our Range of Luxury Rolex Timepieces for Men and Women...
ALL at low prices!
http://scamsite.com/
```

Figure 2.11.3 A sample e-mail showing full headers

the "From:", "To:" and several other headers themselves. As a result, much of the information in these critical headers can be forged – a property which is exploited to good effect by many of those who send unsolicited commercial e-mail (UCE or Spam).

For this reason, it is considered essential to be able to inspect the full headers (i.e. those inserted by the relay hosts) in order to determine the true origins of any e-mail. Unfortunately, most modern e-mail clients hide these headers from end users as they are considered confusing and irrelevant. However, it is usually possible to get the mail software to disclose the headers with a little bit of persuasion. If this cannot be done in the software directly, the headers are almost always present in the file in which received e-mail is stored.

Collecting e-mail

The advent of personal computers in the 1980s, leading to the fulfilment of Pournelle's law, led to a situation where delivery to a local server was no longer sufficient to allow users to work with e-mail efficiently. Either all desktop machines had to be capable of receiving SMTP, which could lead to problems if machines were switched off for prolonged periods, or a mechanism which allowed users to collect their e-mail on demand needed to be devised. In 1984, the first such system was published as RFC918, defining the "Post Office Protocol" as a means for client software to collect e-mail from mailboxes held on a central server. This has evolved into POP-3, which extends the original functionality to support more advanced features.

A further development lies in the design and implementation of IMAP (Internet Message Access Protocol) which, in addition to allowing retrieval and deletion of e-mail on the server, allows users to create and manipulate folders for e-mail storage on the server as well.

From an investigative point of view, the retrieval mechanism has little or no effect on the content of the messages retrieved, but may be significant if we need to examine a mailbox which is still active on a server. If we know which retrieval protocol was in use, we can ensure that checks are carried out for all features supported by that protocol.

Other standards

In addition to these open standards, vendors such as AOL[4] and Microsoft have their own communication standards for mail sending and retrieval. The proprietary, closed, nature of these standards means that little detail about them is available, but messages sent and received through these systems usually have to cross the public Internet at some point in their lives, meaning that they may still carry SMTP headers.

2.11.3 Chat rooms and chat form

Many web site and Internet communities/social network sites offer areas where users can chat to each other in real-time, or near real-time. Since these chat room or forum systems are web based they leave similar traces to web pages and, on the server, there is usually a record of the IP address used to post messages along with time and login information associated with each message. Users can adopt any identity they wish, but ultimately the messages they exchange can be traced back to the originating machines through IP addresses.

2.11.4 Instant messengers and peer-to-peer (P2P) software

Instant messengers, such as AOL messenger, GoogleTalk, MSN messenger, ICQ, IRC and Jabber operate in a manner similar to peer-to-peer (P2P) file sharing software such as eMule, Gnutella, BitTorrent and others. Although generally falling under the heading of P2P, they actually tend to operate in either a client-server mode or a true peer-to-peer mode.

In client-server mode a single “exchange” node receives messages from one user and forwards them to the others in the conversation. In this situa-tion, the client machine may have some logs of who was contacted and, if the examiner is lucky, even a log file containing details of the data exchanged. It is unlikely, however, to contain any IP information other than the address of the server which was acting as the gateway for the conversation.

In true peer-to-peer mode, one or more machines may act as a “telephone directory” (also known as a “Super Node”), listing those users who are currently online and, where file sharing is involved, the resources they are offering out to the rest of the community. Once the client has consulted the directory, direct contact is established with the other party in the conversation (much more like a real telephone call), and both ends of the conversation need to know the other’s IP address in order to exchange data.

In both situations, there is no requirement for the software to record anything at all about the data exhange, but many programs do keep records of recent activity, including IP addresses in use at the time. Again, if

an IP address and/or a username can be found, then it may be possible to carry out tracing through liaison with Internet Service Providers and server owners.

In some systems, the Super Nodes are not fixed points, but are elected periodically by the other peer-to-peer clients on the network, thus they can be difficult to identify and will change over time.

2.11.5 Anonymising techniques

Proxies In many networks, traffic is routinely diverted via a "proxy server" which sits between the local network and the wider Internet. These proxies can help network management by providing two services: filtering out unwanted content (i.e. acting as censors to prevent access to unwanted or illegal material) and acting as large local caches of material that is regularly accessed by a large number of users. By keeping copies of material on the local network, the amount of traffic generated on the Internet can be reduced and local users can be given faster access to the most used material.

From the point of view of a server, a proxy looks like just another client and, apart from some which disclose their nature in request headers, cannot be distinguished from any other client program. As a result, a network of 200 machines operating through a proxy will appear as a single IP address to anything outside the local network.

For network investigators, this poses several challenges, not least because of the existence of *Anonymising Proxies*, which are designed to hide all details of the point of origin of requests sent through the proxy. Requests made via such a proxy can be traced back to the proxy, but rarely, if ever, back to the real origin.

Onion skin routing Proxies act as simple relays, using the protocols al-ready in existence on the Internet for communications, and it is (in theory at least) possible to intercept and moni tor communications which use these protocols. There is another way of creating anonymous communications channels which use encryption to carry any data at all from one point to another, without the intermediaries being able to read the content of any messages.

This method uses a chain of relays, each of which has its own encryp-tion/decryption in place. The sender determines which nodes are going to act as relays for it and builds a pre-determined route. It then encrypts the data to be sent using the keys for the relays, in the right order, to create a series of wrappers or "skins" around the data, and sends the message to the first relay. This relay removes its "skin" to find the next layer of encryption and the address of the next relay to send to.

Each node is, therefore, only aware of the node that sent the message to it, and the node it is sending to. Relay nodes cannot read the message because it has been encrypted with another node's key. Only the target node has the right key to decrypt the final message.

The description of this as "Onion Skin Routing" comes from an analogy with peeling an onion, layer by layer, until the core is exposed. The trans-mitted data are the core, the skins are the encryption and routing addresses. "Pass the parcel" would be an equally good name.

Interception of a message at a node yields nothing more than the previous and next nodes and the encrypted data to be sent to the next node. Identification of origin and target at intermediate nodes is, therefore, impossible without extremely powerful decryption systems.

2.11.6 Web "drops"

A final twist comes from the use of web-based e-mail systems. Normally, any messages sent through these systems are subject to the same rules as any other mail sent using SMTP, although it may be difficult to determine which machine was used to access the web mail server.

In recent times, it has been seen that some criminals/terrorists have started to use webmail in a novel way, exploiting one of the features to exchange messages without sending them. This makes it difficult to show that a message has been sent or received.

Instead of typing in a message and sending it, as would be the normal case, these users are creating messages and saving them in the online system as "drafts" which are stored to be sent at some later date and time. Their accomplices can log on to the same e-mail account, using WWW access, from anywhere in the world, to read and edit messages in the same way. Thus, instant communication between two conspirators on opposite sides of the world can be achieved, without any data being explicitly transmitted from one to the other.

In principle, this is similar to the old spy trick of leaving secret messages in hollow tree stumps or other "dead letter drops" for collection by another agent at some time in the future.

2.12 Mobile phones and PDAs

Personal Digital Assistants started life as replacements for FiloFax-style paper-based personal organisers in the 1990s, but their functionality has largely been supplanted by equivalent functions which are now offered in even "low-end" mobile phones. As a result, although there are still many PDAs in existence, this section will concentrate on evidential opportunities presented by mobile phones. Generally speaking, with the exception of communications data, the material available from PDAs will be similar to that found on phones, and many phones and PDAs are built around the same mobile operating system (e.g. Windows CE/PocketPC/Mobile, Palm OS, Linux & Symbian).

Because these devices are designed to be portable and in-use for prolonged periods, they are battery powered. Older devices used standard replaceable power cells, but modern units are equipped with internal rechargeable batteries which can power the device for several days at a time between recharges.

The sheer number of different physical connections used on these devices over the years, and the different power standards adopted by different manufacturers mean that it is vital for anyone seizing equipment of this type to at least attempt to find the power adapter and interface cables associated with it. This is not always possible, so any good examination facility should have kits similar to that shown in Figure 2.12.1 to allow connections to the majority of devices. Fortunately, the emergence of the USB standard means that more and more manufacturers are electing to use standard cables instead of their older propietary connectors.

We can consider these devices to have three main storage areas (two for PDAs) which may be present: internal memory, internal memory and SIM card memory (not on PDAs).

Figure 2.12.1 A typical set of interface cables found in a mobile device examination kit

Removable memory

Most modern devices have a slot where a memory storage card can be inserted into the device to expand the storage space available and allow data to be exchanged with other devices. These memory cards conform to the same standards as for other devices with removable storage, and usually use the same standard FAT[1] filesystem found in those devices. As a result, the examination of the memory cards can be carried out using standard forensic data examination tools, and the principles and mechanisms described in preceding chapters apply. Care must be taken, however, to ensure that the examiner is aware of how the device's operating system handles the MAC timestamps and file deletion, as portable devices may exhibit unusual behaviours (e.g. not setting access times, or completely erasing data when a file is deleted), leading to different interpretations of their contents and activity.

Fortunately, because these storage devices are designed to be moved between hosts, they do not lose data when power is removed and can be treated as hard discs for examination purposes, although an adapter such as that shown in Figure 2.12.2 may be required to allow the software to read the contents.

Internal memory

The internal memory of a portable device needs to fulfill two functions: it must operate as both the device's primary memory for program execution, and as a filesystem for data storage. Depending on how the device has been programmed, there may be distinct areas within memory set aside for this, or it may move memory from one use to another on demand.

Exactly how this memory is organised may not be visible to an external user and the distinction between RAM and ROM may not be obvious. Memory organisation may be particularly difficult to determine if the device provides an emulation mode, where it appears to be a removable storage device when connected to a general purpose computer, as it may attempt to provide a single unified view of all storage areas as if they were a single device

Figure 2.12.2 A multi-format removable memory card reader connected via a forensic bridge write-blocker

Many devices provide two different modes of operation when connected to a computer: one for synchronisation of personal data (e.g. diaries, e-mail, notes etc.) and one for transfer of data files. It is likely that neither of these modes actually provides a true picture of the contents of internal memory and it may be necessary to either dismantle the device and connect it to specialist diagnostic equipment, or to install a data recovery program on the device.

Installation of the data recovery program, of course, violates ACPO Principle 1, because it entails modifying the state of the device. However, ACPO Principle 2 allows for this to happen when the person per-forming the installation and examination is qualified to provide an explanation of their actions and how they may have affected any recovered evidence.

It should be noted that the internal memory of these devices generally needs to be faster and more flexible than removable storage and tends to use the same type of technology as that found in desktop and laptop computers. When power is lost, the contents of the memory will also be lost, although recent studies have suggested that it may be possible to access data in volatile memory after a period of a few minutes, or even longer if the memory is chilled. This result has not, at the time of writing, been fully validated as an acceptable method for handling volatile storage so should be considered an avenue of last resort.

SIM cards and mobile telephony

[Note: the discussion in this section relates mainly to the GSM (Global System for Mobile Communications) network with some relevance to the UMTS (Universal Mobile Telecommunications System) network. These are the standards across Europe. Other network standards are used in other parts of the world.]

Finally, in the case of mobile phones, we can consider the SIM (Sub-scriber Identity Module) or U-SIM (UMTS Subscriber Identity Module) card (Figure 12.8.3). These cards are part of a general family of "smart" cards [40] or UICCs (Universal Integrated Circuit Cards) which contain small quantities of memory and/or processing capability.

For simplicity, we will concentrate on the SIM card, but USIM provides similar features, albeit in a more advanced way.

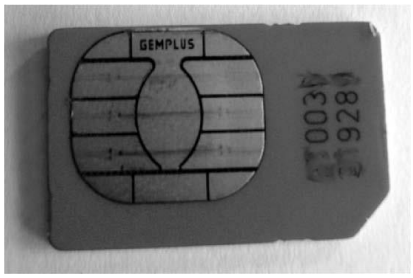

Figure 2.12.3 A standard small-format modern SIM card

As the name suggests, a SIM card provides a way of associating a handset with a subscriber to allow access to the mobile phone network. In order for the network to know where to route call data to, it requires both a logical and physical address (equivalent to the IP and MAC addresses used in the IP suite (see Section 2.12.3). In the mobile phone network, the physical address is the IMSI (International Mobile Subscriber Identity) which is unique to the SIM and programmed in during manufacture. The logical address is the telephone number associated with the SIM through a database which maps the telephone number to the IMSI.

Thus, in order to connect a call, the network matches the telephone number dialled to the IMSI and the call can be routed through to the IMSI's current location.

Another physical identifier exists within the handset, in the form of the IMEI (International Mobile Equipment Identity), and this can be used to identify the handset in use with a particular IMSI or to inform the network that the handset is not permitted to connect if it has been reported stolen.

The SIM standard provides for several small storage areas within the module. Some of the more useful of these, from an investigative perspective, tend to be the following:

Operator data: an area reserved for the network provider to identify which network services and restrictions are in force for the subscriber. For example, this area may list networks which must NOT be connected to, which networks to look for when abroad, whether or not the user is permitted to use the handset abroad etc. One entry which may be of interest in this category is that of the cell in which the handset was last switched off.

Personal data: an area for storage of data created by the SIM user. Typically this may hold their personal address/phone book, short-code dialling numbers, calls made and received and SMS (text) messages sent and received.

PIN/PUK status: this records whether or not a PIN (Personal Identification Number) is required to activate the SIM along with how many times an incorrect PIN has been presented. When the number of incorrect PIN entries exceeds a threshold, the user must enter the Personal Unlock Key (PUK) set by the network operator. Again, there is a fixed num-ber of attempts available for entry of the PUK. If this is exceeded, the

SIM becomes effectively unusable. Examiners should take care to check PIN/PUK status before attempting any other access to the card.

Cell-site analysis

Another fruitful area of investigation involving mobile phones is that of cell-site analysis. This attempts to plot the location where a phone is, or has been used, based on properties of the mobile phone network itself. The GSM network is made up of a network of transmitter/receiver masts, each of which provides coverage for at least one cell. A single mast may operate in all directions at once, or may have directional antennae providing signals in different directions. In rural and/or areas of low population density the masts are usually a few km apart, whereas in urban/highly populated areas the masts need to be much closer together (a few hundred m at most) in order to be able to cope with the number of handsets in the area.

As a handset moves from location to location, it polls the network to determine which mast can provide the best service to it. As a "better" cell is detected, the network switches the handset from the current cell to the new one (the process of "cell handover"), ensuring that calls in progress are not interrupted. Thus, the handset is always aware of which cell it is located in, and the network always knows, approximately, where the handset is based on the antenna in use and the signal strength (which gives a guide to distance from the mast). Cells are typically drawn as idealised hexagons, but in reality the shape of each cell is affected by attenuation of signal strength caused by environmental influences such as trees, hills and buildings.

By mapping signal strengths around a mast, it is possible to predict where cell handover will occur for each network and, if a log of cells used by a handset can be obtained, produce an approximate map of the route taken by the handset.

When a handset is switched off, it signals this fact to the network and both the handset and network record the cell in which it was last active.

When calls are made, the network produces a Call Detail Record (CDR) for billing purposes. Usually this contains, at least, the starting and ending cells for the call along with called number and duration. Although not necessarily a complete list of all cells used during the call, it may be possible to estimate the movement of the handset by plotting signal strengths and predicting cell handovers.

The handset's constant interaction with the network is, possibly, the biggest threat to continuity of evidence, as any interaction post-discovery may result in alteration of the device's contents.

The ACPO Good Practice Guide for Computer Based Electronic Evidence contains detailed advice on the handling of mobile phones. Use of Faraday cages/bags is strongly recommended for the seizure, transport and examination of these devices, but it should be remembered that if a phone cannot contact a cell, it will continue to broadcast using stronger and stronger signals until it either finds a cell or its power supply is exhausted. Since the networks use different frequencies (GSM uses 900 MHz, 1800 MHz, 950 MHz, 1900 MHz, 400 MHz and 450 MHz depending on coun-try), it should not be assumed that all Faraday cages/bags will isolate the handset from all networks. Different Faraday cages/bags will have different levels of permeability for the various frequencies.

2.13 GPS

GPS, the Global Positioning System, uses a network of satellites in orbit around the planet. Each satellite has a unique identity and a well-defined orbital path combined with an accurate clock. Each satellite broadcasts information about the current time and its orbit.

By identifying the satellites "visible" from any point on the planet and performing a calculation based on the data sent by the satellites, a GPS receiver can calculate exactly where it is above the surface of the planet.

Although designed originally for military use, GPS technology has reached the consumer market, where a typical receiver costs around the

Figure 2.13.1 A typical consumer GPS unit for use in-car

Same as a PDA. The use of this technology has become very popular with drivers, walkers, sailors and amateur pilots as an aid to navigation, so it is becoming increasingly common for computer crime units to be asked to examine these devices for possible evidence.

GPS devices (Figure 12.3.1) are showing signs of convergence with PDA technology, with some GPS units now having media player facilities, the ability to connect to mobile phones via Bluetooth, WiFi potential, and general-purpose storage functions as well.

Examination of GPS equipment can be more problematic than exam1-ination of phones. Some systems use versions of the operating systems designed for PDAs, while others are based on entirely proprietary navigation software.

If interrogation of the device can be carried out, though, it may be possible to retrieve navigation data about the point of last use, lists of favourite places and pre-planned routes, in addition to other us-age information which depends on the other features available on the device.

Figure 12.3.2 Personal media players

2.13 Other personal technology

Most people now own at least a few other examples of personal technology in the form of digital cameras, media players etc. On the whole, these devices are single-function but can operate as external storage devices when connected to more powerful computers. Some, however, are appearing with the convergent properties described above. In addition to their main functions, they offer PDA-type functionality, WiFi connectivity etc. with all the evidential opportunites and investigative problems that those features create.

Unit III: Introduction to Computer Forensics, Use of Computer Forensics in Law Enforcement, Computer Forensics Assistance to Human Resources / Employment Proceedings, Computer Forensics Services, Benefits of Professional Forensics Methodology, Steps Taken by Computer Forensics Specialists, Who Can Use Computer Forensic Evidence?, Case Histories, Case Studies.

3.1 INTRODUCTION TO COMPUTER FORENSICS

Computer forensics, also referred to as computer forensic analysis, electronic discovery, electronic evidence discovery, digital discovery, data recovery, data discovery, computer analysis, and computer examination, is the process of methodically examining computer media (hard disks, diskettes, tapes, etc.) for evidence. A thorough analysis by a skilled examiner can result in the reconstruction of the activities of a computer user.

In other words, computer forensics is the collection, preservation, analysis, and presentation of computer-related evidence. Computer evidence can be useful in criminal cases, civil disputes, and human resources/employment proceedings.

Far more information is retained on a computer than most people realize. It's also more difficult to completely remove information than is generally thought. For these reasons (and many more), computer forensics can often find evidence of, or even completely recover, lost or deleted information, even if the information was intentionally deleted.

Computer forensics, although employing some of the same skills and software as data recovery, is a much more complex undertaking. In data recovery, the goal is to retrieve the lost data. In computer forensics, the goal is to retrieve the data and interpret as much information about it as possible.

The continuing technological revolution in communications and information exchange has created an entirely new form of crime: cyber crime or computer crime. Computer crime has forced the computer and law enforcement professions to develop new areas of expertise and avenues of collecting and analyzing evidence. This is what has developed into the science of computer forensics. The process of acquiring, examining, and applying digital evidence is crucial to the success of prosecuting a cyber criminal. With the continuous evolution of technology, it is difficult for law enforcement and computer professionals to stay one step ahead of technologically savvy criminals. To effectively combat cyber crime, greater emphasis must be placed in the computer forensic field of study, including but not limited to financial support, international guidelines and laws, and training of the professionals involved in the process, as well as the following subject matter:

Computer crime

- The computer forensic objective
- The computer forensic priority
- The accuracy versus speed conflict
- The need for computer forensics
- The double tier approach
- Requirements for the double tier approach
- The computer forensics specialist

Computer Crime

According to industry analysts, there are currently 657 million people online world-wide. That figure is expected to rise to 794 million by 2009. This represents a lot of data interchange. Unfortunately many small businesses, and even large organizations, do not know how to properly protect their sensitive data, thus leaving the door open to criminals.

Computers can be involved in a wide variety of crimes including white-collar crimes, violent crimes such as murder and terrorism, counterintelligence, economic espionage, counterfeiting, and drug dealing. A 2003 FBI survey reported that the average bank robbery netted $6,900, whereas the average computer crime netted $900,000 [1]. The Internet has made targets much more accessible, and the risks involved for the criminal are much lower than with traditional crimes. A per-son can sit in the comfort of his home or a remote site and hack into a bank and transfer millions of dollars to a fictitious account, in essence robbing the bank, without the threat of being gunned down while escaping. One hears of such technological crimes almost daily, thus creating a perception of lawlessness in the cyber world. The same FBI survey revealed that both public and private agencies face se-rious threats from external as well as internal sources. Out of the 849 organizations that responded to the survey, 30% claimed theft of proprietary information, 23% reported sabotage of data or their networks, 35% experienced system penetration from an outside source, and 12% claimed financial fraud. More alarming is the ease of access to sensitive data employees have within the organization. Fifty-nine percent of the organizations involved in the survey reported employees having unauthorized access to corporate information.

Recently a survey was conducted to determine where the FBI was focusing their computer forensic efforts. An alarming 74% of their workload is centered on white-collar crime. This type of crime includes health care fraud, government fraud including erroneous IRS and Social Security benefit payments, and financial institution fraud. These are high-dollar crimes made easy by technology. The other 26% of the workload is split equally among violent crime (child pornography, interstate theft), organized crime (drug dealing, criminal enterprise), and counter-terrorism and national security. As shown by this survey, computer crime is widespread and has infiltrated areas unimaginable just a few years ago. The FBI caseload has gone from near zero in 1985 to nearly 10,000 cases in 2003. It is no doubt considerably higher today. They have gone from two part-time scientists to 899 personnel in regional field offices throughout the country. Technology has brought this field of study to the forefront.

Roles of a Computer in a Crime

A computer can play one of three roles in a computer crime. A computer can be the target of the crime, it can be the instrument of the crime, or it can serve as an evidence repository storing valuable information about the crime. In some cases, the computer can have multiple roles. It can be the “smoking gun” serving as the instrument of the crime. It can also serve as a file cabinet storing critical evidence. For example, a hacker may use the computer as the tool to break into another computer and steal files, then store them on the computer. When investigating a case, it is important to know what roles the computer played in the crime and then tailor the investigative process to that particular role.

Applying information about how the computer was used in the crime also helps when searching the system for evidence. If the computer was used to hack into a network password file, the investigator will know to look for password cracking software and password files. If the computer was the target of the crime, such as an intrusion, audit logs and unfamiliar programs should be checked. Knowing how the computer was used will help narrow down the evidence collection process. With the size of hard drives these days, it can take a very long time to check and analyze every piece of data a computer contains. Often law enforcement

officials need the information quickly, and having a general idea of what to look for will speed the evidence collection process.

The Computer Forensic Objective

The objective in computer forensics is quite straightforward. It is to recover, analyze, and present computer-based material in such a way that it is useable as evidence in a court of law. The key phrase here is useable as evidence in a court of law. It is essential that none of the equipment or procedures used during the examination of the computer obviate this.

The Computer Forensic Priority

Computer forensics is concerned primarily with forensic procedures, rules of evidence, and legal processes. It is only secondarily concerned with computers. Therefore, in contrast to all other areas of computing, where speed is the main concern, in computer forensics the absolute priority is accuracy. One talks of completing work as efficiently as possible, that is, as fast as possible without sacrificing accuracy.

Accuracy versus Speed

In this seemingly frenetic world where the precious resource of time is usually at a premium, pressure is heaped upon you to work as fast as possible. Working under such pressure to achieve deadlines may induce people to take shortcuts in order to save time.

In computer forensics, as in any branch of forensic science, the emphasis must be on evidential integrity and security. In observing this priority, every forensic practitioner must adhere to stringent guidelines. Such guidelines do not encompass the taking of shortcuts, and the forensic practitioner accepts that the precious resource of time must be expended in order to maintain the highest standards of work.

The Computer Forensics Specialist

A computer forensics specialist is the person responsible for doing computer forensics. The computer forensics specialist will take several careful steps to identify and attempt to retrieve possible evidence that may exist on a subject computer system:

1. Protect the subject computer system during the forensic examination from any possible alteration, damage, data corruption, or virus introduction.
2. Discover all files on the subject system. This includes existing normal files, deleted yet remaining files, hidden files, password-protected files, and encrypted files.
3. Recover all (or as much as possible) of discovered deleted files.
4. Reveal (to the extent possible) the contents of hidden files as well as temporary or swap files used by both the application programs and the operating system.
5. Accesses (if possible and if legally appropriate) the contents of protected or encrypted files.
6. Analyze all possibly relevant data found in special (and typically inaccessible) areas of a disk. This includes but is not limited to what is called unallocated space on a disk (currently unused, but possibly the repository of previous data that is relevant evidence), as well as slack space in a file (the remnant area at the end of a file, in the last assigned disk cluster, that is unused by cur-rent file data but once again may be a possible site for previously created and relevant evidence).

7. Print out an overall analysis of the subject computer system, as well as a listing of all possibly relevant files and discovered file data. Further, provide an opinion of the system layout; the file structures discovered; any discovered data and authorship information; any attempts to hide, delete, protect, or encrypt information; and anything else that has been discovered and appears to be relevant to the overall computer system examination.
8. Provide expert consultation and/or testimony, as required.

Who Can Use Computer Forensic Evidence?

Many types of criminal and civil proceedings can and do make use of evidence revealed by computer forensics specialists.

- Criminal Prosecutors use computer evidence in a variety of crimes where in-criminating documents can be found: homicides, financial fraud, drug and embezzlement record-keeping, and child pornography.
- Civil litigations can readily make use of personal and business records found on computer systems that bear on fraud, divorce, discrimination, and harassment cases. Insurance companies may be able to mitigate costs by using discovered computer evidence of possible fraud in accident, arson, and workman's compensation cases.
- Corporations often hire computer forensics specialists to find evidence relating to sexual harassment, embezzlement, theft or misappropriation of trade secrets, and other internal/confidential information.
- Law enforcement officials frequently require assistance in pre-search warrant preparations and post-seizure handling of the computer equipment. The use of computer forensics in law enforcement is discussed in detail in the next section and throughout the book.
- Individuals sometimes hire computer forensics specialists in support of possible claims of wrongful termination, sexual harassment, or age discrimination.

3.2 USE OF COMPUTER FORENSICS IN LAW ENFORCEMENT

If there is a computer on the premises of a crime scene, the chances are very good that there is valuable evidence on that computer. If the computer and its contents are examined (even if very briefly) by anyone other than a trained and experienced computer forensics specialist, the usefulness and credibility of that evidence will be tainted.

- **Choosing a Computer Forensics Specialist for a Criminal Case**

When you require the services of a computer forensics specialist, don't be afraid to shop around. There are an increasing number of people who claim to be experts in the field. Look very carefully at the level of experience of the individuals involved. There is far more to proper computer forensic analysis than the ability to retrieve data, especially when a criminal case is involved. Think about computer forensics just as you would any other forensic science and look for a corresponding level of expertise.

The bottom line is that you will be retaining the services of an individual who will likely be called to testify in court to explain what he or she did to the computer and its data. The court will want to know that individual's own level of training and experience, not the experience of his or her employer. Make sure you find someone who not only has the expertise and experience, but also the ability to stand up to the scrutiny and pressure of cross-examination.

3.3 COMPUTER FORENSICS ASSISTANCE TO HUMAN RESOURCES/EMPLOYMENT PROCEEDINGS

Computer forensics analysis is becoming increasingly useful to businesses. Computers can contain evidence in many types of human resources proceedings, including sexual harassment suits, allegations of discrimination, and wrongful termination claims. Evidence can be found in electronic mail systems, on network servers, and on individual employee's computers. However, due to the ease with which computer data can be manipulated, if the search and analysis is not performed by a trained computer forensics specialist, it could likely be thrown out of court.

Employer Safeguard Program

As computers become more prevalent in businesses, employers must safeguard critical business information. An unfortunate concern today is the possibility that data could be damaged, destroyed, or misappropriated by a discontented individual.

Before an individual is informed of their termination, a computer forensic specialist should come on-site and create an exact duplicate of the data on the individual's computer. In this way, should the employee choose to do anything to that data before leaving, the employer is protected. Damaged or deleted data can be re-placed, and evidence can be recovered to show what occurred. This method can also be used to bolster an employer's case by showing the removal of proprietary in-formation or to protect the employer from false charges made by the employee.

Whether you are looking for evidence in a criminal prosecution or civil suit or determining exactly what an employee has been up to, you should be equipped to find and interpret the clues that have been left behind. This includes situations where files have been deleted, disks have been reformatted, or other steps have been taken to conceal or destroy the evidence. For example, did you know

- What Web sites have been visited
- What files have been downloaded
- When files were last accessed
- Of attempts to conceal or destroy evidence
- Of attempts to fabricate evidence
- That the electronic copy of a document can contain text that was removed from the final printed version
- That some fax machines can contain exact duplicates of the last several hundred pages received
- That faxes sent or received via computer may remain on the computer indefinitely That email is rapidly becoming the communications medium of choice for businesses
- That people tend to write things in email that they would never consider writing in a memorandum or letter
- That email has been used successfully in criminal cases as well as in civil litigation that email is often backed up on tapes that are generally kept for months or years that many people keep their financial records, including investments, on computers.

3.4 COMPUTER FORENSICS SERVICES

No matter how careful they are, when people attempt to steal electronic information (everything from customer databases to blueprints), they leave behind traces of their activities. Likewise, when people try to destroy incriminating evidence contained on a computer (from harassing memos to stolen technology), they

leave behind vital clues. In both cases, those traces can prove to be the smoking gun that successfully wins a court case. Thus, computer data evidence is quickly becoming a reliable and essential form of evidence that should not be overlooked.

A computer forensics professional does more than turn on a computer, make a directory listing, and search through files. Your forensics professionals should be able to successfully perform complex evidence recovery procedures with the skill and expertise that lends credibility to your case. For example, they should be able to perform the following services:

- Data seizure
- Data duplication and preservation
- Data recovery
- Document searches
- Media conversion
- Expert witness services
- Computer evidence service options
- Other miscellaneous services

Data Seizure

Federal rules of civil procedure let a party or their representative inspect and copy designated documents or data compilations that may contain evidence. Your computer forensics experts, following federal guidelines, should act as this representative, using their knowledge of data storage technologies to track down evidence. Your experts should also be able to assist officials during the equipment seizure process. "

Data Duplication and Preservation

When one party must seize data from another, two concerns must be addressed: the data must not be altered in any way, and the seizure must not put an undue burden on the responding party. Your computer forensics experts should acknowledge both of these concerns by making an exact duplicate of the needed data. Because duplication is fast, the responding party can quickly resume its normal business functions, and, because your experts work on the duplicated data, the integrity of the original data is maintained.

Data Recovery

Using proprietary tools, your computer forensics experts should be able to safely recover and analyze otherwise inaccessible evidence. The ability to recover lost evidence is made possible by the expert's advanced understanding of storage technologies. For example, when a user deletes an email, traces of that message may still exist on the storage device. Although the message is inaccessible to the user, your experts should be able to recover it and locate relevant evidence.

Document Searches

Your computer forensics experts should also be able to search over 200,000 electronic documents in seconds rather than hours. The speed and efficiency of these searches make the discovery process less complicated and less intrusive to all parties involved.

Media Conversion

Some clients need to obtain and investigate computer data stored on old and un-readable devices. Your computer forensics experts should extract the relevant data from these devices, convert it into readable formats, and place it onto new storage media for analysis.

Expert Witness Services

Computer forensics experts should be able to explain complex technical processes in an easy-to-understand fashion. This should help judges and juries comprehend how computer evidence is found, what it consists of, and how it is relevant to a specific situation (see sidebar, "Provide Expert Consultation and Expert Witness Services").

PROVIDE EXPERT CONSULTATION AND EXPERT WITNESS SERVICES COMPUTERS

Expert Testimony

- Has testified multiple times as an expert witness in computers and computer forensics in circuit court
- Regularly testify as an expert witness in computers and computer forensics in federal court for U.S. attorney's offices

Computer Expertise

- Belongs to the Computer Crime Investigators Association
- Trained in the forensic examination of computers (PC & Mac), having con-ducted examinations in countless cases including child exploitation, homicide, militia, software piracy, and fraud
- Has testified in state and federal courts as an expert in computers, computer forensics, the Internet, and America Online; often as an expert witness for U.S. attorney's offices
- Is thoroughly familiar with both computer hardware and software, having writ-ten software and repaired and assembled computers
- Teaches computer crime investigation, including computer search and seizure, for the Institute of Police Technology and Management
- Regularly consults with law enforcement officers in the search and seizure of computers
- Has provided forensic training to numerous law enforcement officers and corporate security officers
- Regularly consulted by other forensic examiners for advice in difficult cases

Training Given as Expert in Computer Crimes

- Law Enforcement and Corrections Technology Symposium and Exhibition
- Bureau of Justice Statistics/Justice Research Statistics Association.

Computer Evidence Service Options

Your computer forensics experts should offer various levels of service, each de-signed to suit your individual investigative needs. For example, they should be able to offer the following services:

- Standard service
- On-site service
- Emergency service
- Priority service
- Weekend service

Standard Service

Your computer forensics experts should be able to work on your case during nor-mal business hours until your critical electronic evidence is found. They must be able to provide clean rooms and ensure that all warranties on your equipment will still be valid following their services.

On-Site Service

Your computer forensics experts should be able to travel to your location to per-form complete computer evidence services. While on-site, the experts should quickly be able to produce exact duplicates of the data storage media in question. Their services should then be performed on the duplicate, minimizing the disrup-tion to business and the computer system. Your experts should also be able to help federal marshals seize computer data and be very familiar with the Federal Guide-lines for Searching and Seizing Computers.

Emergency Service

After receiving the computer storage media, your computer forensics experts should be able to give your case the highest priority in their laboratories. They should be able to work on it without interruption until your evidence objectives are met.

Priority Service

Dedicated computer forensics experts should be able to work on your case during normal business hours (8:00 A.M. to 5:00 P.M., Monday through Friday) until the evidence is found. Priority service typically cuts your turnaround time in half.

Weekend Service

Computer forensics experts should be able to work from 8:00 A.M. to 5:00 P.M., Saturday and Sunday, to locate the needed electronic evidence and will continue working on your case until your evidence objectives are met. Weekend service depends on the availability of computer forensics experts.

Other Miscellaneous Services

Computer forensics experts should also be able to provide extended services. These services include

- Analysis of computers and data in criminal investigations
- On-site seizure of computer data in criminal investigations
- Analysis of computers and data in civil litigation.
- On-site seizure of computer data in civil litigation
- Analysis of company computers to determine employee activity
- Assistance in preparing electronic discovery requests
- Reporting in a comprehensive and readily understandable manner
- Court-recognized computer expert witness testimony
- Computer forensics on both PC and Mac platforms
- Fast turnaround time

Recover Data You Thought Was Lost Forever

Computers systems may crash. Files may be accidentally deleted. Disks may acci-dentally be reformatted. Computer viruses may corrupt files. Files may be acciden-tally overwritten. Disgruntled employees may try to destroy your files. All of these can lead to the loss of your critical data. You may think it's lost forever, but com-puter forensics experts should be able to employ the latest tools and techniques to recover your data.

In many instances, the data cannot be found using the limited software tools available to most users. The advanced tools that computer forensics experts utilize allow them to find your files and restore them for your use. In those

instances where the files have been irreparably damaged, the experts' computer forensics expertise allows them to recover even the smallest remaining fragments.

Advise You on How to Keep Your Data and Information Safe from

Theft or Accidental Loss

Business today relies on computers. Your sensitive client records or trade secrets are vulnerable to intentional attacks from, for example, computer hackers, disgruntled employees, viruses, and corporate espionage. Equally threatening, but far less considered, are unintentional data losses caused by accidental deletion, computer hardware and software crashes, and accidental modification.

Computer forensics experts should advise you on how to safeguard your data by such methods as encryption and back-up. The experts can also thoroughly clean sensitive data from any computer system you plan on eliminating.

Your files, records, and conversations are just as vital to protect as your data. Computer forensics experts should survey your business and provide guidance for improving the security of your information. This includes possible information leaks such as cordless telephones, cellular telephones, trash, employees, and answering machines.

3.5 BENEFITS OF PROFESSIONAL FORENSICS METHODOLOGY

The impartial computer forensics expert who helps during discovery will typically have experience on a wide range of computer hardware and software. It is always beneficial when your case involves hardware and software with which this expert is directly familiar, but fundamental computer design and software implementation is often quite similar from one system to another. Experience in one application or operating system area is often easily transferable to a new system.

Unlike paper evidence, computer evidence can often exist in many forms, with earlier versions still accessible on a computer disk. Knowing the possibility of their existence, even alternate formats of the same data can be discovered. The discovery process can be served well by a knowledgeable expert identifying more possibilities that can be requested as possibly relevant evidence. In addition, during on-site premises inspections, for cases where computer disks are not actually seized or forensically copied, the forensics expert can more quickly identify places to look, signs to look for, and additional information sources for relevant evidence. These may take the form of earlier versions of data files (memos, spreadsheets) that still exist on the computer's disk or on backup media or differently formatted versions of data, either created or treated by other application programs (word processing, spreadsheet, email, timeline, scheduling, or graphic).

Protection of evidence is critical. A knowledgeable computer forensics professional should ensure that a subject computer system is carefully handled to ensure that

- No possible evidence is damaged, destroyed, or otherwise compromised by the procedures used to investigate the computer.
- No possible computer virus is introduced to a subject computer during the analysis process
- Extracted and possibly relevant evidence is properly handled and protected from later mechanical or electromagnetic damage
- A continuing chain of custody is established and maintained Business operations are affected for a limited amount of time, if at all
- Any client-attorney information that is inadvertently acquired during a forensic exploration is ethically and legally respected and not divulged.

Steps Taken by Computer Forensics Specialists

The computer forensics specialist needs to complete an Evidence Identification and Retrieval Checklist. He or she should take several careful steps to identify and attempt to retrieve possible evidence that may exist on a subject's computer system.

3.6 WHO CAN USE COMPUTER FORENSIC EVIDENCE?

Many types of criminal and civil proceedings can and do make use of evidence re-vealed by computer forensics specialists. These are as follows:

- Criminal prosecutors use computer evidence in a variety of crimes where in-criminating documents can be found, including homicides, financial fraud, drug and embezzlement record-keeping, and child pornography.
- Civil litigations can readily make use of personal and business records found on computer systems that bear on fraud, divorce, discrimination, and harassment cases.
- Insurance companies may be able to mitigate costs by using discovered computer evidence of possible fraud in accident, arson, and workman's compensation cases.
- Corporations often hire computer forensics specialists to find evidence relating to sexual harassment, embezzlement, and theft or misappropriation of trade secrets, and other internal and confidential information.
- Law enforcement officials frequently require assistance in pre-search warrant preparations and post seizure handling of the computer equipment. Individuals sometimes hire computer forensics specialists in support of possible claims of wrongful termination, sexual harassment, or age discrimination.

However, there are concerns and problems with computer forensic evidence. Let's examine some of those problems.

Legal Tests

The rules vary from legislation to legislation, but one can give a broad outline of what happens in those countries with a common law tradition—the U.K., U.S., and the so-called old Commonwealth. The law makes distinctions between real evidence, testimonial evidence, and hearsay. Real evidence is that which comes from an inanimate object that can be examined by the court. Testimonial evidence is that which a live witness has seen and upon which he or she can be cross-examined. The hearsay rule operates to exclude assertions made other than those made by the witness who is testifying as evidence of the truth of what is being asserted. The pure hearsay rule is extremely restrictive and has been extensively modified by various statutory provisions. Thus, there are rules about the proving of documents and business books. Bankers' books have separate legislation. Some of the rules apply explicitly to computers, but many do not, although they can be (and have been) interpreted to cover many situations in which computers are involved.

For example, in the U.K. there have been situations where legal rules presumably designed to help the court may in fact hinder it. In practice, these issues may be circumvented. For instance, in a criminal case, evidence may be obtained by in-admissible methods. This evidence, however, then points investigators to admissible sources of evidence for the same sets of circumstances. An example of this could occur during a fraud investigation. In other words, computer search methods are often used to identify allegedly fraudulent transactions, but the evidential items eventually presented in court are paper-based invoices, contract notes, dockets, or other documents. In this manner, the prosecution can demonstrate to the jury the deception or breach of the Companies Act or other specific fraudulent act. Again, in civil litigation the parties may decide to jointly accept computer-based evidence (or not to challenge it) and instead concentrate on the more substantive elements in the dispute. A defendant may prefer to have a substantive defense rather than a technical one based on inadmissibility. Or, again, the legal team may not feel sufficiently competent to embark on a technical challenge.

In the U.S., many practical problems exist around the actual seizure of computers containing evidence. Law enforcement officers must comply with the Fourth Amendment to the U.S. Constitution.

Subject Matter of Computer Forensics

The subject matter of computer forensics can, thus, not be solely concerned with procedures and methods of handling computers, the hardware from which they are made up, and the files they contain. The ultimate aim of forensic investigation is its use in legal proceedings. At the same time, an obsession with common law and judicial rules is likely to inhibit many investigations. It might be a mistake for inquiries not to be commenced simply because of fear of possible inadmissibility. Furthermore, as we have already seen, a number of computer-investigatory methods may turn out not to be directly admissible but may nevertheless be useful in locating noncomputer evidence that is admissible.

One may have to take a somewhat pragmatic view of the precise bounds of the subject matter, but it should still be possible to define its core activities. It might help to explore the way in which forensic science in general has developed and then see what expectations one might reasonably have of computer forensics.

Although forensic science was already well established, and indeed forms a cen-tral feature of many of Conan Doyle's Sherlock Holmes stories published from 1892 onwards, up until the 1970s, each forensic scientist tended to develop his or her own methods and present them ad hoc to juries. Obviously, reliance was placed on descriptions of methods used by others, but for courts, the tests of whether to believe the forensic evidence were the manner of presentation—the supposed eminence of the forensic scientist and the skill of the opposition lawyer (or rival expert who might be called). During the 1970s, a more formal checklist-based approach was introduced. This was partly to bring about standardization as between different laboratories and partly in response to the criticism (in the U.K.) that arose over such controversial cases as the Birmingham Six. In the U.K. Home Office Forensic Service, these checklists were devised by senior staff. Obviously, such checklists are revised in the light of experience—the publication of new specialist research or ad-verse experience during a trial. An increasingly used feature of modern practice is quality control, which involves work being checked by an other wise uninvolved coworker before being offered to external scrutiny. In any event, the broad tests for evidence include

Authenticity: Does the material come from where it purports?

Reliability: Can the substance of the story the material tells be believed and is it consistent? In the case of computer-derived material, are there reasons for doubting the correct working of the computer?

Completeness: Is the story that the material purports to tell complete? Are there other stories that the material also tells that might have a bearing on the legal dispute or hearing?

Freedom from interference and contamination: Are these levels acceptable as a result of forensic investigation and other post event handling.

Divergences from Conventional Forensic Investigation

There will be divergences from the expectations of more traditional areas of foren-sic investigation. The main reason is the rate of change of computer technology. The devisor of a test for the presence of a prohibited drug, an explosive, fabric fibers, bodily tissues, and the like, can expect that over a period of time, the test may be improved or shown to be defective, but, the need for the test and most of its es-sential details will probably not change. However, in computers, newness and obsolesce is the norm.

For example, a key feature of computer forensics is the examination of data media: new forms and methods of data storage occur at intervals of less than 4 years. The floppy disk of 13 years ago was in 5.25 inch format and held 360 k. The current equivalent is 3.5 inches and holds 1.44 MB, and much higher densities are expected soon. A typical hard-disk size on a PC of the same date was 20-30 MB, was in 5.25 inch form, and used modified frequency modulation (MFM) controller technology. Today most PCs have hard disks in excess of 1750 MB in 2.5 inch or even

1.5 inch form using integrated development environment (IDE) or run length limited (RLL) technology. On minis and mainframes, data may be held on redundant array of independent (or inexpensive) disks (RAID), where individual files may be split and spread over eight or more separate disk surfaces. Similar changes have taken place in tape technology and the use of erasable programmable read-only memory (EPROMs).

Computer architectures have gone through profound changes in the same short period. PCs have become much more powerful, the large central mainframe is now a rarity, and large companies are now served by a multiplicity of smaller computers that all interact via a complex network.

Computer peripherals keep changing as well. Modems and network routers have become intelligent, and digitizing scanners are fairly common devices. They can be subverted, for example, for forgery.

Wide-area telecom methods are being used more and more. These provide op-portunities for both high-tech criminals and forensic investigators. The protocols they use also keep changing.

The foregoing simply lists technological changes. Similar changes have taken place in computer applications; these, in turn, have affected the type of information one might expect to find held in a computer. For example, over the same 13 years, the following technological changes have taken place:

- The growth of email, both locally within large organizations and worldwide. The growth of client/server applications.
- The software outcome of the more complex hardware architectures. The client/server situation (software on).
- The ability of a PC or small local machine to interact with software and data held on other nonlocal machines and large mainframes in a way that appears to be seamless to the user. One key effect of this is that a computer document often does not exist in some computer equivalent of a filing cabinet, but, rather, is assembled on demand by the activity of one computer drawing information from many others.
- The evidence of a transaction or event may, therefore, only be provable by the presentation of all the records from all the computers involved, plus an expla-nation of how the assembly of the report relied on took place.
- The greater use of Electronic Data Interchanges (EDIs) and other forms of computer-based orders, bills of lading, payment authorizations, etc. EDIs have very complex structures, with some evidence being held in computers owned by the counter-parties and some by the EDI supplier/regulator.
- Computer graphics: computer-aided design (CAD) methods, particularly those that provide an element of auto completion or filling-in of basic design ideas. More extended, easier-to-use databases.
- The greater use of computer-controlled procedures (sales, dispatch, and emergency services; computer-controlled processes; traffic control; and manufacturing).
- The methods of writing and developing software. There is much greater use of libraries of procedures (of new computer language models). For example, object-oriented programming environments and new, more formal methods of program development; standards and methods of testing have also changed [5].

As a result, computer forensic methods may not have the time in which to establish themselves, or the longevity, that more traditional chemistry- and physics-based forensics enjoy. Nevertheless, the usual way in which specific forensic methods become accepted is via publication in a specialist academic journal. For example, a forensic scientist seeking to justify a methodology in court can do so by stating that it is based on a specific published method that had not up to the point of the hearing been criticized.

CASE HISTORIES

One of the fundamental principles of computer investigation is the need to follow established and tested procedures meticulously and methodically throughout the investigation. At no point of the investigation is this more critical than at the stage of initial evidence capture. Reproducibility of evidence is the key. Without the firm base of solid procedures, which have been strictly applied, any subsequent antire-pudiation attempts in court will be suspect, and the case as a whole will likely be weakened.

There have been several high-profile cases recently where apparently solid cases have been weakened or thrown out on the basis of inappropriate consideration given to the integrity and reproducibility of the computer evidence. This may hap-pen for several reasons. Lack of training is a prime culprit. If the individuals in-volved have not been trained to the required standards, or have received no training at all, then tainted or damaged computer evidence is the sad but inevitable result.

Another frequent cause is lack of experience. Not only lack of site experience, but also inappropriate experience of the type of systems, might be encountered. One of the most difficult on-site skills is knowing when to call for help. It is essential that a sympathetic working environment is created such that peer pressure or fear of loss of status and respect does not override the need to call for help. Easier said than done, perhaps, but no less essential for that reason.

Finally, sloppiness, time pressure, pressure applied on-site, fatigue, and carelessness have all been contributory factors in transforming solid computer evidence into a dubious collection of files. These totally avoidable issues are related to individual mental discipline, management control and policy, and selecting appropri-ate staff to carry out the work. There are issues with which one cannot sympathize. This is bad work, plain and simple.

Ultimately, any time the collection of computer evidence is called into ques-tion, it is damaging to everyone who is a computer forensic practitioner; it is in everyone's best interest to ensure that the highest standards are maintained.

To use a rather worn phrase from an old American police series (*Hill Street Blues*): "Let's be careful out there!"

Taken for a Ride

A sad, but all too frequent story, from prospective clients: I've just spent $15,000 on a Web site and got taken for a ride. I cannot find the con man now and all I have is an alias and a pay-as-you-go mobile number. Can you help me please?

What Can You Do?

It is strongly recommended that people dealing with entities on the Internet need to make sure they know who they are dealing with before they enter into any transaction or agreement. If you cannot obtain a real-world address (preferably within the juris-diction in which you live), then think twice about going any further. Always question the use of mobile phone numbers—they should set alarm bells ringing! This task is made easier in the U.K., as all mobile numbers start with 077xx, 078xx, or 079xx. Pagers start with 076xx. From April 28, 2001, on, all old mobile, pager (those that do not begin 07), special rate, and premium rate numbers stopped working.

If you do want to proceed with the transaction, then use a credit card rather than a debit card or other type of money transfer; then at least you will have some protection and only be liable for $50 rather than having your entire bank account cleaned out. In terms of tracing a suspect like the one in the preceding, your com-puter forensic experts should be able to trace emails around the world; and, by act-ing quickly and in conjunction with legal firms, they should be able to track individuals down to their homes. An application for a civil search order can then allow entry and the experts will be able to secure all electronic evidence quickly and efficiently. Internet cafés are sometimes more of a problem, but it is remarkable how many users go to the trouble of trying to disguise their tracks only to end up sitting in exactly the same seat every time they visit the same Café. So, yes, your computer forensic experts can help, but by taking the proper precautions, you would not need to call them in the first place.

Abuse of Power and Position

This message is by no means new; in fact, it could be said that it has been repeated so many times in so many forums that it is amazing that management still falls foul of the following circumstances. In recent months, investigators at Vogon

International Limited have been asked to examine computer data for evidence of fraud. On one occasion, the client was a charity, and on the second, a multinational company.

In both cases, fraud, totaling hundreds of thousands of dollars was uncovered. The modus operandi of the suspects was very similar in both cases. Bogus compa-nies were set up and invoices were submitted for payment. The fraudsters were in a position to authorize the payment of the invoices and had the power to prevent unwelcome scrutiny of the accounts.

In addition, one of the fraudsters was paying another member of the staff to turn a blind eye to what was happening. On further investigation, this member of the staff was obviously living beyond his means.

The message is simple: whether you are a multinational company or a small business, the possibility of fraud is ever present. While not wishing to fuel paranoia, traditional checks and balances must be in place to ensure that those trusted members of the staff who have power cannot abuse their positions.

Secure Erasure

Now, let's touch on this "old chestnut" again, because it appears to be the source of considerable confusion and misinformation. Vogon's customer base seems to be polarized into two main camps: those who desperately want to retain their data and fail, often spectacularly, to do so and those who wish to irrevocably destroy their data, and frequently fail in a similarly dramatic manner.

The latter may be criminals who wish to cover their tracks from the police or legitimate business organizations who wish to protect themselves from confidential information falling into the wrong hands. Fundamentally, the issues are the same. The legitimate destruction of data is ultimately a matter of management responsibility, which requires a considered risk analysis to be carried out.

To the question, Can data be securely erased?, the answer is, self-evidently, yes. If you were to ask, Is it straightforward or certain?, it depends, would be the answer.

Many systems are in use for securely erasing data from a wide range of media. Some are effective, some completely ineffective, and some partially effective. It is the latter situation that causes concern and, frequently, not an inconsiderable amount of embarrassment.

Those systems that absolutely destroy data do so in a manner that is total, un-equivocal, and final; there can exist no doubt as to their effectiveness. Systems that are sold as being completely effective but that are fundamentally flawed are obviously flawed. With only cursory analysis, this is evident, so these are (or should be) swiftly disregarded.

Vogon is regularly asked to verify the destruction of data by many of their large clients. What they find is that frequently only a fraction of a sample sent is correctly or accurately deleted. RAID systems are a prime candidate for chaos. Certain revisions of drive firmware can present special challenges; in some cases, even the software used defeats the eraser. The list of such software is long and growing.

Vogon is often asked for advice on this issue. The answer is always the same. If the destruction of data has more value than the drive, physically destroy the drive. Crushing is good; melting in a furnace is better. If the drive has more value than the data, what are you worrying about?

CASE STUDIES

Over the years, Vogon's data-recovery laboratories have seen pretty much every-thing that can happen to a computer, no matter how incredible, whether it is a ge-ologist who, in testing for minerals, inadvertently blew up his own laptop, or the factory worker who covered the computer running the production line in maple syrup. The list is now so long that the incredible has become almost mundane. For-tuitously, two in the latest of a long line of incredible recoveries recently occurred, so, it seemed appropriate to include them as case studies.

Case Study One: The Case of the Flying Laptop

Picture the scene: police rushing into premises on the ninth floor of a building. Almost immediately thereafter, a laptop accelerates rapidly ground ward out of the window of the aforementioned premises.

As long ago as 1687, Sir Isaac Newton predicted with uncanny accuracy the in-evitable conclusion to this action: namely, the laptop (or to be strictly accurate, large number of pieces of a former laptop) coming to rest with a singular lack of grace on the ground. Luckily, no one was injured by the impact. The resultant bag of smashed laptop components arrived at Vogon's laboratory for a forensically sound data recovery.

The laptop computer had impacted the ground across its front edge at an angle, forcing the hard disk drive assembly to go completely through the screen of the lap-top. The highly delicate spatial relationship between heads, flexures, platters, and spindle had become disturbed, and the bed of the drive unit was not concave. This imparted an oscillation in two dimensions during drive operation. The drive electronics were destroyed in the impact. After an evening's work by a highly skilled hardware engineer, it was determined that a full fix was possible, and a perfect image was taken. Vogon had no knowledge of whether the chap was guilty, but they bet he was in shock when the evidence was presented.

Case Study Two: The Case of the Burned Tapes

This case does not involve true forensic investigation, but it does highlight the fact that it is important never to give up on a job, no matter how seemingly hopeless it appears.

Sets of digital audio tape (DAT) tapes were sent to Vogon from a loss adjuster. The DAT tapes were caught in a fire, which had engulfed a company's head office and wiped out the primary trading infrastructure. The company's IT systems had been at the center of the blaze, and this had unfortunately raised the magnetic media on the surface of the servers hard drives past its curie point. The DAT tapes had, rather inadvisably as it turned out, not been stored off-site. They were, however, stored a little way from the center of the blaze.

Despite this, the DAT tapes arrived in a rather sorry condition. The plastic casing had melted to, around, and onto the tapes, and the whole mechanism was fused into a homologous glob. It is fair to say the tapes were sent to Vogon with the full expectation that they would be declared unrecoverable and used as the basis from which to make a loss settlement.

This recovery involved hours of work from both hardware and tape recovery engineers. The tapes were carefully cut away from the molten mass and treated for fire damage. The next stage was to rehouse the tapes and pass them forward to the tape recovery team. Following a number of complex stages, the recovery team was able to extract a stream of data from the tapes that accounted for some 95% of the original data stored on the company's tape backups.

The result was a company up and running in a matter of days rather than weeks, or, more likely, never. It also resulted in a significant reduction in the claims settlement by the loss adjuster and business continuity for the unfortunate company.

Unit IV: Types of Military Computer Forensic Technology, Types of Law Enforcement: Computer Forensic Technology, Types of Business Computer Forensic Technology, Specialized Forensics Techniques, Hidden Data and How to Find It, Spyware and Adware, Encryption Methods and Vulnerabilities, Protecting Data from Being Compromised, Internet Tracing Methods.

4.1 TYPES OF MILITARY COMPUTER FORENSIC TECHNOLOGY

The U.S. Department of Defense (DoD) cyber forensics includes evaluation and in-depth examination of data related to both the trans- and post cyberattack periods. Key objectives of cyber forensics include rapid discovery of evidence, estimation of potential impact of the malicious activity on the victim, and assessment of the in-tent and identity of the perpetrator. Real-time tracking of potentially malicious activity is especially difficult when the pertinent information has been intentionally hidden, destroyed, or modified in order to elude discovery. The information directorate's cyber forensic concepts are new and untested. The directorate entered into a partnership with the National Institute of Justice via the auspices of the National Law Enforcement and Corrections Technology Center (NLECTC) located in Rome, New York, to test these new ideas and prototype tools. The Computer Forensics Experiment 2000 (CFX-2000) resulted from this partnership. This first-of-a-kind event represents a new paradigm for transitioning cyber forensic technology from military research and development (R&D) laboratories into the hands of law enforcement. The experiment used a realistic cyber crime scenario specifically designed to exercise and show the value added of the directorate-developed cyber forensic technology.

The U.S. Department of Defense (DoD) cyber forensics includes evaluation and in-depth examination of data related to both the trans- and post cyber attack periods. Key objectives of cyber forensics include rapid discovery of evidence, estimation of potential impact of the malicious activity on the victim, and assessment of the in-tent and identity of the perpetrator. Real-time tracking of potentially malicious activity is especially difficult when the pertinent information has been intentionally hidden, destroyed, or modified in order to elude discovery. The information directorate's cyber forensic concepts are new and untested. The directorate entered into a partnership with the National Institute of Justice via the auspices of the National Law Enforcement and Corrections Technology Center (NLECTC) located in Rome, New York, to test these new ideas and prototype tools. The Computer Forensics Experiment 2000 (CFX-2000) resulted from this partnership. This first-of-a-kind event represents a new paradigm for transitioning cyber forensic technology from military research and development (R&D) laboratories into the hands of law enforcement. The experiment used a realistic cyber crime scenario specifically designed to exercise and show the value added of the directorate-developed cyber forensic technology.

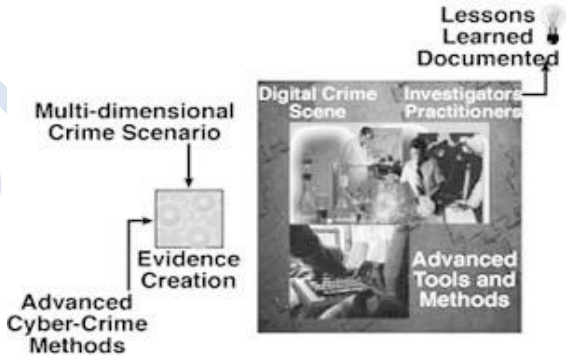

FIGURE 4.1 CFX-2000 schematic (© 2002, Associated Business Publications. All rights reserved).

Digital evidence, perform case management and time lining of digital events, auto-mate event link analysis, and perform steganography detection. The results of CFX-2000 verified that the hypothesis was largely correct and that it is possible to ascertain the intent and identity of cyber criminals. As electronic technology continues its explosive growth, researchers need to continue vigorous R&D of cyber forensic technology in preparation for the onslaught of cyber reconnaissance probes and attacks.

4.2 TYPES OF LAW ENFORCEMENT COMPUTER FORENSIC TECHNOLOGY

As previously defined, computer forensics involves the preservation, identification, extraction, and documentation of computer evidence stored in the form of magnetically encoded information (data). Often the computer evidence was created transparently by the computer's operating system and without the knowledge of the computer operator. Such information may actually be hidden from view and, thus, special forensic software tools and techniques are required to preserve, identify, ex-tract, and document the related computer evidence.

Computer forensics tools and techniques have proven to be a valuable resource for law enforcement in the identification of leads and in the processing of computer-related evidence. Computer forensics tools and techniques have become important resources for use in internal investigations, civil lawsuits, and computer security risk management.

Forensic software tools and methods can be used to identify passwords, logons, and other information that is automatically dumped from the computer memory as a transparent operation of today's popular personal computer operating systems. Such computer forensic software tools can also be used to identify backdated files and to tie a diskette to the computer that created it.

Law enforcement and military agencies have been involved in processing computer evidence for years. This section touches very briefly on issues dealing with Windows NT®, Windows® 2000, XP and 2003 and their use within law enforcement computer forensic technology.

Windows XP and Windows 2003 are operating systems that are often used on notebook and desktop computers in corporations and government agencies. Thus, they are currently the operating systems most likely to be encountered in computer investigations and computer security reviews. Be advised that this chapter does not cover the use of *black box* computer forensics software tools. Those tools are good for some basic investigation tasks, but they do not offer a full computer forensics solution. Furthermore, such approaches are all but useless in computer security risk assessments. Such assessments usually require that searches and file listings be con-ducted overtly or even covertly from a single floppy diskette.

Computer Evidence Processing Procedures

Processing procedures and methodologies should conform to federal computer evidence processing standards. Computer processing procedures have also been developed for the U.S. Treasury Department.

Training and certification programs have also been developed for the International Association of Computer Investigation Specialists (IACIS). For these reasons, computer forensic trainers and instructors should be well qualified to teach the correct computer-processing methods and procedures.

Preservation of Evidence

Computer evidence is fragile and susceptible to alteration or erasure by any number of occurrences. Computer forensic instructors should expose their trainees to bit stream backup theories that ensure the preservation of all storage levels that may contain evidence. For example, Safe Back software overcomes some of the evidence weaknesses inherent in black box computer forensics approaches (see sidebar, "Mirror Image Backup Software"). Safe Back technology can be purchased from New Technologies, Inc. and has become a worldwide standard in making mirror image backups since 1990, when it was developed based on requirements then established by the U.S. Treasury Department and the IACIS.

MIRROR IMAGE BACKUP SOFTWARE

SafeBack is used to create mirror-image (bit-stream) backup files of hard disks or to make a mirror-image copy of an entire hard disk drive or partition. The process is analogous to photography and the creation of a photo negative. Once the photo neg-ative has been made, several exact reproductions can be made of the original. Unlike a photo, SafeBack image files cannot be altered or modified to alter the reproduc-tion. This is because SafeBack is an industry standard self-authenticating computer forensics tool that is used to create evidence-grade backups of hard drives.

With the release of SafeBack version 3.0 or higher, the integrity of SafeBack files is maintained through the use of two separate mathematical hashing processes that rely upon the National Institute of Standards and Technology (NIST)-tested Secure Hash Algorithm256 (SHA256). Users of prior versions of SafeBack are encouraged to upgrade to take advantage of the greater levels of accuracy achieved with version 3.0. Information about upgrades can be found on the Internet at *http://www.forensics-intl.com/*. The upgrade of SafeBack has new and added features and it takes into account the last sector error finding by NIST concerning the older SafeBack version 2.0.

Backup image files created with SafeBack can be written to any writable magnetic storage device, including small computer system interface (SCSI) tape backup units. SafeBack preserves all the data on a backed-up or copied hard disk, including inactive or deleted data. Backup image files can be restored to another system's hard disk. Remote operation via a parallel port connection allows the hard disk on a remote PC to be read or written by the master system. A date- and time-stamped audit trail maintains a record of SafeBack operations during a session, and when the default is used an SHA256 hash is recorded in the output audit file. This hash can be used to cross-validate the accuracy of the process with any other software utility that relies upon the NIST-tested SHA256 algorithm. To avoid possible claims that the SafeBack image file may have been altered after the fact, SafeBack now safeguards the internally stored SHA256 values. Any alterations of computer data are quickly brought to the attention of the operator of the program when the SafeBack image file is restored.

Simply put, SafeBack is a DOS-based utility used to back up and restore hard disks. SafeBack picks up every last bit of data-unused and erased data included—on the original disk and stores it in a tape or disk file (or series of files). SafeBack can take that same backup file and recreate the original disk on your own system. SafeBack does not write or otherwise modify the original system and can (and should) be started from a boot diskette.

SafeBack also has a couple of *derivative* operating modes. The first is Verify mode, where restoring from a backup disk is done, but the data is thrown away. This is more useful than it first appears to be because it allows the operator of the pro-gram to scan his or her backups to make sure that they will read back without errors, without having to go through the setup required by a standard Safe Back restore procedure. The other derivative operation is Copy, which feeds the Restore function directly with the output of the Backup function, with no intermediate files. This is less useful than it first appears to be. If the operator of Safe Back is considering making a copy, he might as well make a backup image file and then restore it as needed.

PRIMARY USES

The primary uses of Safe Back are as follows:

- Used to create evidence-grade backups of hard disk drives on Intel-based computer systems.
- Used to exactly restore archived Safe Back images to another computer hard disk drive of equal or larger storage capacity.
- Used as an evidence preservation tool in law enforcement and civil litigation matters.
- Used as an intelligence gathering tool by military agencies.

Trojan Horse Programs

The need to preserve the computer evidence before processing a computer should be clearly demonstrated by the computer forensic instructor through the use of programs designed to destroy data and modify the operating systems. The partici-pant should be able to demonstrate his or her ability to avoid destructive programs and traps that can be planted by computer users bent on destroying data and evidence. Such programs can also be used to covertly capture sensitive information, passwords, and network logons.

Computer Forensics Documentation

The documentation of forensic processing methodologies and findings is important. This is even true concerning computer security risk assessments and internal audits, because without proper documentation, it is difficult to present findings. If the security or audit findings become the object of a lawsuit or a criminal investigation, then

documentation becomes even more important. Thus, the computer forensic instructor should also teach the participant the ins and outs of computer evidence processing methodology (which facilitates good evidence-processing documentation and good evidence chain of custody procedures). The benefits will be obvious to investigators, but they will also become clear to internal auditors and computer security specialists.

File Slack

The occurrence of random memory dumps in hidden storage areas should be discussed and covered in detail during workshops. Techniques and automated tools that are used to capture and evaluate file slack should be demonstrated in a training course. Such data is the source of potential security leaks regarding passwords, network logons, email, database entries, and word processing documents. These security and evidence issues should also be discussed and demonstrated during the training course. The participants should be able to demonstrate their ability to deal with slack and should demonstrate proficiency in searching file slack, documenting their findings, and eliminating the security risk.

Data-Hiding Techniques

Trade secret information and other sensitive data can easily be secreted using any number of techniques. It is possible to hide diskettes within diskettes and to hide entire computer hard disk drive partitions. These issues should be discussed in any computer forensics training course from a detection standpoint, as well as from a security risk standpoint. Tools that help in the identification of such anomalies should be demonstrated and discussed (like AnaDisk™ [see sidebar, "AnaDisk Diskette Analysis Tool"]) in the training course. Participants should be required to demonstrate their understanding of such issues. This aspect of the training becomes especially important during the last day of the course when the participants are called on to extract their Certificate of Completion from a *special* floppy diskette. Data-hiding courses are only open to classified government agencies and businesses that have a demonstrated need to know about this kind of information as outlined in a company's training policies. This is because the information covered in a data-hiding course can be used to defeat government computer security review processes and techniques.

ANADISK DISKETTE ANALYSIS TOOL

AnaDisk turns your PC into a sophisticated diskette analysis tool. The software was originally created to meet the needs of the U.S. Treasury Department in 1991. It is primarily used to identify data storage anomalies on floppy diskettes and generic hardware in the form of floppy disk controllers; bios are needed when using this software. It works at a very low level and makes maximum use of the floppy diskette hardware. The software also has limited search capabilities and can be used to copy abnormal diskettes. It can also be used to write data at a physical sector level and to format diskettes using any number of combinations.

AnaDisk has the capability to duplicate floppy diskettes, but this feature is used primarily with odd diskette formats (in cases like the FBI Russian mole case of sus-pected spy Robert Phillip Hanssen). However, standard duplication of floppy diskettes is more easily accomplished with New Technology Inc.'s CopyQM (see sidebar "CopyQM: Diskette Duplication Software"), which has been upgraded and certified by the U.S. DoD for making copies of diskettes used in classified computer security risk reviews.

In other words, AnaDisk can be used to analyze floppy diskettes when doing computer evidence consulting work, which involves abnormal floppy diskettes or data storage issues tied to floppy diskettes. It can also be used in data-hiding courses to create data-hiding areas by adding extra sectors and tracks to floppy diskettes and in writing data to unformatted floppy diskettes.

PRIMARY USES

- Security reviews of floppy diskettes for storage anomalies
- Duplication of diskettes that are nonstandard or that involve storage anomalies Editing diskettes at a physical sector level
- Searching for data on floppy diskettes in traditional and nontraditional storage areas

- Formatting diskettes in nontraditional ways for training purposes and to illustrate data hiding techniques

COPYQM: DISKETTE DUPLICATION SOFTWARE

CopyQM Plus essentially turns a personal computer into a diskette duplicator. In a single pass, diskettes are formatted, copied, and verified. This capability is useful for computer forensics specialists and computer security specialists who need to pre-configure floppy diskettes for specific uses and duplicate them.

Classified government agencies and government contractors are required to perform periodic examinations of government computer systems to determine if classified data may reside on nonclassified computer systems. Programs like New Technology Inc.'s Text Search Plus (see sidebar, "Text Search Plus") and Text Search NT were designed specifically for this purpose and they are both certified by the U.S. DoD for use in classified government risk assessments. CopyQM is also certified by the U.S. DoD for use in the duplication of "search disks" used in classified U.S. government computer risk reviews.

CopyQM Plus can also create self-extracting executable programs that can be used to duplicate specific diskettes. This feature makes CopyQM an ideal tool for use in security reviews because once a CopyQM disk-creation program has been created, it can be used by anyone to create pre-configured security risk assessment diskettes. When the resulting program is run, the diskette image of the original diskette will be restored on multiple diskettes automatically.

This process requires little technical knowledge and it allows computer specialists to delegate more of the security risk assessment responsibilities to employees with minimal technical knowledge. The diskette images can also be password protected when the diskette images are converted to self-extracting programs. This is helpful when you want to keep computer forensic and computer security software tools away from curious hands.

PRIMARY USES

- The program is used to archive the image of a floppy diskette and to create one or more duplicate copies of the master diskette when desired.
- It can be used to make one or more copies of all areas of a *normal* floppy diskette. Thus, it basically turns your PC into a diskette duplication machine. This can be helpful when you need to repeatedly make copies of diskettes for training classes. It is also helpful for making multiple copies of evidence diskettes.
- It can be used to automatically create and serialize software stored on floppy diskettes. This type of *branding* can be helpful in creating copies of diskettes that will be shared among several lawyers or with the court.
- CopyQM Plus can be used to password protect the contents of an entire floppy diskette. This is helpful when diskettes are shared over the Internet and when security is a concern.
- CopyQM Plus can be used to create virus-scanned floppy diskette tool kits con-figured for repeated tasks performed by computer forensics specialists, electronic data personnel (EDP), auditors and computer security specialists. The software is particularly helpful in creating computer incident response tool kit diskettes.
- CopyQM Plus can be used to send a normal diskette over the Internet.

E-Commerce Investigations

A new Internet *forensic tool* has recently been introduced that aims to help educators, police, and other law enforcement officials trace the past World Wide Web activity of computer users. Net Threat Analyzer™, from Gresham, Oregon-based New Technology Inc. (NTI), can be used to identify past Internet browsing and email activity done through specific computers. The software analyzes a computer's disk drives and other storage areas that are generally unknown to or beyond the reach of most general computer users.

Kids can figure out ways to prevent their parents from finding anything on their machine, but Net Threat Analyzer goes back in after the fact where things are easier to detect. New Technology Inc. has made its Net Threat Analyzer avail-able free of charge to computer crime specialists, school officials, and police.

The program is booted from a floppy disk and uses filtering tools to collect data on users' basic browsing and email history. It flags possible threats, such as any-thing dealing with drugs, bombs, country codes, or pornography. Web sites change so often that it's difficult to keep up with which ones are porn or drug sites.

For example, http://www.whitehouse.gov, is the official White House Web site, and www.whitehouse.com is a pornography site. If Junior's been to whitehouse.com 500 to 700 times, it will make it through most net nanny software, but it will raise a red flag with the Net Threat Analyzer.

The software was designed to help prevent situations like the tragedies at Columbine High School in Littleton, Colorado, and Thurston High School in Springfield, Oregon, where weapons were made by teenagers who had downloaded the instructions from the Internet.

New Technology Inc., which specializes in computer forensics tools and training, has posted order forms for its software on its Web site at http://www.forensics-intl.com. The tool is not available to the public, but a special version can be purchased by Fortune 500 companies, government agencies, military agencies, and consultants who have a legitimate need for the software.

Dual-Purpose Programs

Programs can be designed to perform multiple processes and tasks at the same time. They can also be designed for delayed tasking. These concepts should be demonstrated to the training participants during the course through the use of specialized software. The participant should also have hands-on experience with these programs.

Text Search Techniques

New Technology Inc. has also developed specialized search techniques and tools that can be used to find targeted strings of text in files, file slack, unallocated file space, and Windows swap files. Each participant will leave their training class with a licensed copy of their TextSearch Plus™ software and the necessary knowledge to conduct computer security reviews and computer related investigations (see sidebar, "Text Search Plus").

TEXT SEARCH PLUS

Text Search Plus was specifically designed and enhanced for speed and accuracy in security reviews. It is widely used by classified government agencies and corporations that support these agencies. The software is also used by hundreds of law enforcement agencies throughout the world in computer crime investigations.

This software is used to quickly search hard disk drives, zip disks, and floppy diskettes for key words or specific patterns of text. It operates at either a logical or physical level at the option of the user. Text Search Plus has been specifically designed to meet the requirements of the government for use in computer security exit reviews from classified government facilities. The current version is approximately 25% faster than prior versions. It is also compatible with FAT 12, FAT 16, and FAT 32 DOS-based systems. As a result, it can be used on Windows 98, 2000, XP, and 2003 systems. Tests indicate that this tool finds more text strings than any other forensic search tool. It is sold separately and is also included in several of the New Technology Inc. tool suites. As a stand alone tool, it is ideal for security risk assessments. When security spills are identified, they can easily be eliminated with New Technology Inc.'s M-Sweep™ program.

PRIMARY USES

- Used to find occurrences of words or strings of text in data stored in files, slack, and unallocated file space
- Used in exit reviews of computer storage media from classified facilities
- Used to identify data leakage of classified information on nonclassified computer systems

- Used in internal audits to identify violations of corporate policy
- Used by Fortune 500 corporations, government contractors, and government agencies in security reviews and security risk assessments
- Used in corporate due diligence efforts regarding proposed mergers
- Used to find occurrences of keywords strings of text in data found at a physical sector level
- Used to find evidence in corporate, civil, and criminal investigations that involve computer-related evidence
- Used to find embedded text in formatted word processing documents (Word Perfect™ and fragments of such documents in ambient data storage areas).

Fuzzy Logic Tools Used to Identify Unknown Text

New Technology Inc. has also developed a methodology and tools that aid in the identification of relevant evidence and unknown strings of text. Traditional computer evidence searches require that the computer specialist know what is being searched for. However, many times not all is known about what may be stored on a given computer system. In such cases, fuzzy logic tools can provide valuable leads as to how the subject computer was used. The training participants should be able to fully understand these methods and techniques. They should also be able to demonstrate their ability to use them to identify leads in file slack, unallocated file space, and Windows swap files. Each training participant should also leave the class with a licensed copy of New Technology Inc.'s Filter_G™ software (see sidebar, "Intelligent Forensic Filter").

INTELLIGENT FORENSIC FILTER

This forensic filter utility is used to quickly make sense of nonsense in the analysis of ambient data sources (Windows swap/page files, file slack, and data associated with erased files). Filter_G is a unique fuzzy logic filter that was awarded patent number 6,345,283 by the U.S. Patent Office. It is used to quickly identify patterns of English language grammar in ambient data files. Such an analysis can be helpful in making quick assessments about how a specific com-puter was used and the nature of prior English language communications that were involved in the past uses of a subject computer. The program can be used as a sampling tool and it is particularly useful when used to evaluate Windows swap/page files.

Be aware that the functionality of this software was contained in New Technology Inc.'s Filter_I prior to March, 2003. Since that time the functionality was substantially enhanced and incorporated into this program as a stand-alone utility.

PRIMARY USES

- Used as an intelligence gathering tool for quick assessments of a Windows swap/page file to identify past communications on a targeted computer
- Used as a data sampling tool in law enforcement, military, and corporate investigations.
- Used to quickly identify patterns of English language grammar in ambient data sources.
- Used to identify English language communications in erased file space.

Disk Structure

Participants should be able to leave a training course with a good understanding of how computer hard disks and floppy diskettes are structured and how computer evidence can reside at various levels within the structure of the disk. They should also demonstrate their knowledge of how to modify the structure and hide data in obscure places on floppy diskettes and hard disk drives.

Data Encryption

A computer forensics course should cover, in general, how data is encrypted; it should also illustrate the differences between good encryption and bad encryption. Furthermore, demonstrations of password-recovery software should be

given regarding encrypted WordPerfect, Excel, Lotus, Microsoft Word, and PKZIP files. The participant should become familiar with the use of software to *crack* security associated with these different file structures.

Matching a Diskette to a Computer

New Technology Inc. has also developed specialized techniques and tools that make it possible to conclusively tie a diskette to a computer that was used to create or edit files stored on it. Unlike some *special* government agencies, New Technology Inc. re-lies on logical rather than physical data storage areas to demonstrate this technique. Each participant is taught how to use special software tools to complete this process.

Data Compression

The participant should be shown how compression works and how compression programs can be used to hide and disguise sensitive data. Furthermore, the participant should learn how password-protected compressed files can be broken; this should be covered in hands-on workshops during the training course.

Erased Files

The training participant should be shown how previously erased files can be recov-ered by using DOS programs and by manually using data-recovery techniques. These techniques should also be demonstrated by the participant, and cluster chaining will become familiar to the participant.

Internet Abuse Identification and Detection

The participant should be shown how to use specialized software to identify how a targeted computer has been used on the Internet. This process will focus on computer forensics issues tied to data that the computer user probably doesn't realize exists (file slack, unallocated file space, and Windows swap files).

The Boot Process and Memory Resident Programs

The participant should be able to take part in a graphic demonstration of how the operating system can be modified to change data and destroy data at the whim of the per-son who configured the system. Such a technique could be used to covertly capture keyboard activity from corporate executives, for example. For this reason, it is important that the participants understand these potential risks and how to identify them.

4.3 TYPES OF BUSINESS COMPUTER FORENSIC TECHNOLOGY

Finally, let's briefly look at the following types of business computer forensics technology:

- Remote monitoring of target computers
- Creating trackable electronic documents
- Theft recovery software for laptops and PCs
- Basic forensic tools and techniques
- Forensic services available

Remote Monitoring of Target Computers

Data Interception by Remote Transmission (DIRT) from Codex Data Systems (CDS), Inc. is a powerful remote control monitoring tool that allows stealth monitoring of all activity on one or more target computers simultaneously from a remote command center. No physical access is necessary. Application also allows agents to remotely seize and secure digital evidence prior to physically entering suspect premises.

Creating Trackable Electronic Documents

There are so many powerful intrusion detection tools that allow the user to create trackable electronic documents that it is beyond the scope of this chapter to mention them all. See "Intrusion Detection Systems" in Chapter 3 for a detailed explanation of some of these tools.

In general, most of these tools identify (including their location) unauthorized intruders who access, download, and view these *tagged* documents. The tools also allow security personnel to trace the chain of custody and chain of command of all who possess the stolen electronic documents.

Theft Recovery Software for Laptops and PCs

If your PC or laptop is stolen, is it smart enough to tell you where it is? According to a recent FBI report, 98% of stolen computers are never recovered. According to Safeware Insurance, 1,201,000 PCs and laptops were stolen in 2002 and 2003, costing owners $7.8 billion dollars [9]. According to a recent joint Computer Security Institute/FBI survey, 72% of the Fortune 1000 companies experienced laptop theft [9].

Nationwide losses to computer component theft cost corporate America over $11 billion a year. So if your company experiences computer-related thefts and you do nothing to correct the problem, there is a 92% chance you will be hit again.

What Is the Real Cost of a Stolen Laptop or PC?

When you lose your wallet, the last thing you think of is how much it is going to cost to replace your wallet. The same is true when equipment (especially a com-puter) is stolen.

Our mothers always told us, an ounce of prevention is worth a pound of cure. They were right. Think about what it really costs to replace a stolen computer.

- The price of the replacement hardware.
- The price of replacing the software.
- The cost of recreating data. If possible at all, do you keep perfect back-ups? The cost of lost production time or instruction time.
- The loss of customer goodwill (lost faxes, delayed correspondence or billings, problems answering questions and accessing data).
- The cost of reporting and investigating the theft, filing police reports and insurance claims.
- The cost of increased insurance.
- The cost of processing and ordering replacements, cutting a check, and the like. If a thief is ever caught, the cost of time involved in prosecution.

PC Phone Home Scenario

Imagine you're a small business owner and you just went out and invested several thousand dollars in the latest laptop computer, fully loaded with all the bells and whistles. On your first business trip with your new computer you leave it (you think) safely hidden in your hotel room while you entertain a client. Later that night you return to your room and your laptop, with all your work in it, has disappeared without a trace. The financial loss is bad enough, but the hours of work you've lost is worse, and the sensitivity of the information in your laptop, if it gets into the hands of the wrong people, could be a disaster.

Now imagine a simple, inexpensive software system that offers real hope of tracking your laptop and pinpointing its location anywhere in the world. Is this re-ally possible or is it just another fanciful hi-tech gimmick from the imagination of the writers of the latest James Bond movie? It's no gimmick. It's PC Phone Home, the latest in computer theft recovery software.

PC Phone Home is a software application that, when installed in your laptop or desktop computer, secretly transmits an electronic message to an email address of your choice. This allows you to track and locate your computer, thus providing the potential for its ultimate recovery as well as apprehension of the thief.

How Does PC PhoneHome Work?

It's simple. First, you install PC PhoneHome on your computer, configuring it to send its recovery information to an email address of your choosing. PC PhoneHome sends a stealth email to your designated email address once a day, or every time you connect to the Internet and are assigned an IP address different from your previous IP address. If your computer is lost or stolen, you report the loss to the police and continue to monitor (with the additional help of the PC PhoneHome Recovery Cen-ter) your designated email address. When your stolen computer accesses the Internet by any method, your lost or stolen computer will send you its stealth email message, informing you of its location.

If you are a registered user of PC PhoneHome, you may seek the PC Phone-Home technical service center's assistance in locating your computer's exact coor-dinates and alerting the local police to recover it. As a side benefit, any other items of your property (like expensive jewelry) that might have been taken at the same time may also be recovered.

A success story, PC PhoneHome has been enthusiastically embraced by police forces, insurance companies, and the computer industry. The product is a natural fit for the security monitoring and Internet service provider (ISP) industry. PC Phone-Home is compatible with all Windows and Macintosh operating systems.

Basic Forensic Tools and Techniques

Today, many computer forensics workshops have been created to familiarize investigators and security personnel with the basic techniques and tools necessary for a successful investigation of Internet and computer-related crimes. So many work-shops have been created that it is beyond the scope of this chapter to mention them all. However, throughout the book, a number of them will be mentioned in detail. Workshop topics normally include: types of computer crime, cyber law basics, tracing email to its source, digital evidence acquisition, cracking passwords, monitor-ing computers remotely, tracking online activity, finding and recovering hidden and deleted data, locating stolen computers, creating trackable files, identifying software pirates, and so on.

Forensic Services Available

Through computer forensic evidence acquisition services, forensic experts for companies like Capitol Digital Document Solutions can provide management with a potent arsenal of digital tools at its disposal. They have the necessary software and hardware to travel to designated sites throughout the world to acquire an exact image of hard drives, tapes, etc. This image is an exact duplication of the source media and allows evaluation within their laboratories with minimal disruption to others. Services include but are not limited to

- Lost password and file recovery
- Location and retrieval of deleted and hidden files
- File and email decryption
- Email supervision and authentication
- Threatening email traced to source
- Identification of Internet activity
- Computer usage policy and supervision
- Remote PC and network monitoring
- Tracking and location of stolen electronic files Honeypot sting operations.
- Location and identity of unauthorized software users Theft recovery software for laptops and PCs.
- Investigative and security software creation
- Protection from hackers and viruses

VIRUS/TROJAN/WORM PROTECTION

The following are tips to avoid a computer virus and trojan and worm programs:

- Don't open attachments sent with junk email or emails from persons you do not know.
- Don't open any files attached to an email if the subject line is questionable or unexpected. If you need to open the file, always save it to your hard drive before doing so.
- Disable the Windows Scripting Host. Recently, Microsoft introduced a "macro" programming language into the core of Windows and IE browsers, which allows Visual Basic scripts to run without the need for specialist soft-ware. Although this can make your computer easier to use (being able to pro-gram shortcuts, or use third-party Visual Basic scripts), it is also a security risk. A typical PC does not need Windows Scripting Host to function correctly, and it should be safe to disable it. To remove Windows Scripting Host from your computer, open up your Control Panel and select the Add/Remove Programs icon. Select the Windows Setup tab, double-click the Accessories section and untick the box next to "Windows Scripting Host."
- Always download files from well-known established and trusted sites. Always know the source of any attachment and file. Install an anti-virus program. This program will scan any file or attachment you get over the Net for known viruses. The program will warn you if any virus is detected.
- Correctly configure the anti-virus software so that it performs as designed. An improperly configured anti-virus software can be as good as no software.
- If your anti-virus program has an automatic virus scanning feature, keep this feature activated just in case you forget to scan for the virus.
- Back up your files on a regular basis. If a virus destroys your files, at least you can replace them with your backup copy. You should store your backup copy in a separate location from your work files, one that is preferably not on your computer.
- Get immediate protection. Configure your anti-virus software to boot automatically on start-up and run at all times. In case you forget to boot up your anti-virus software, configuring it to start by itself will ensure you are always protected.
- Educate yourself about the latest viruses. Many Web sites keep a list of all old and new viruses.
- Write-protect all system and software diskettes using the write-protect tab to prevent any virus from spreading. Using the write-protect tab will prevent viruses from being transmitted to those sensitive or critical system disks.
- Don't boot from a floppy disk. Floppies are a common way viruses are trans-mitted. If you use a floppy while working on your computer, remove it when you shut the machine off or the computer will automatically try to boot from the floppy, perhaps launching any viruses on the disk.
- Don't share floppies. Even a well-meaning friend may unknowingly pass along a virus, trojan horse, or worm.
- Scan floppies before using them. This is always important, but especially if you are using the disk to carry information between one computer and another. You could easily pick up a virus from an insecure network and introduce it into your system. Running a virus scan before launching any of the programs on the disk will prevent infection.
- If you are in a network environment and you get infected, report the virus to your systems administrator, who can them determine the source of the infection. This will ensure that the virus will not spread.

- If you use Microsoft Outlook (not Express) for email, make sure that the Automatic preview feature is disabled. You can find this option under the View menu. Outlook and Outlook Express users are the most targeted. Consider using a different email program.
- If you must use one of the Outlooks, download all of Microsoft's security patches. Add your own address to your Outlook address book, so if it starts sending out messages on its own, at least you'll know about it.
- Keep in mind that often computer viruses are spread by accident, so files you get from a friend who would not purposely infect you still might give you a virus

4.4 SPECIALIZED FORENSICS TECHNIQUES

Threats to the strategic value of your business almost always involve a computer or network because that is where your company's proprietary information and business processes are located. A simple and virtually undetectable fraud that posts a few cents to a phony account can reap a perpetrator thousands of dollars flowing through accounts payable. A malicious change to an individual's personnel records could cost the person a job and a career. Divulging a company's financial records could damage it on Wall Street, in the marketplace, and before shareholders. Cor-porate espionage can steal trade secrets. Posting libelous information on the Inter-net about a company or individual can damage a reputation beyond recovery. Employees of a company might be stealing from it or using company resources to work for themselves, or they can be using excessive work time to surf pornographic sites and play games.

Computer forensics investigators examine computer hardware and software using legal procedures to obtain evidence that proves or disproves allegations. Gathering legal evidence is difficult and requires trained specialists who know computers, the rules of evidence gathering, and how to work with law enforcement authorities.

Computer forensics examiners should be called in when a threat to a company's business and reputation is serious. Any organization that does not have a way to de-tect and stop malicious behavior can be victimized with no legal recourse. Preserv-ing evidence according to Federal Rules of Evidence gives choices that otherwise would not exist. When an intruder attacks or steals from an organization, the abil-ity or threat to get law enforcement involved may be the only way to reduce the dam-age or prevent future occurrences. Gathering computer evidence is also useful for confirming or dispelling concerns about whether an illegal incident has occurred and for documenting computer and network vulnerabilities after an incident.

Companies employ computer forensics when there is serious risk of informa-tion being compromised, a potential loss of competitive capability, a threat of law-suits, or potential damage to reputation and brand. Some companies regularly use forensic investigations to check employee computers. In theory, employees are less tempted to stray when they know they are being watched.

On the other hand, when the cost of a forensic investigation exceeds potential gain, there is little reason to use it. Companies have used legal evidence gathering to drive home points with employees and external intruders even though the cost of investigations exceeded recovery. Usually, however, a full-scale investigation is not needed to stop an inappropriate action such as surfing that wastes time. Computer forensics also may not be needed when computers and networks play a minor role in an incident or threat, but this may not always be clear. The relationship between the computer and an event under inquiry is critical, and sometimes until a forensics examination has been done, one cannot know whether a computer was a significant part of an event or not.

Legal Evidence

A computer forensics examiner always should gather and preserve evidence according to Federal Rules of Evidence. The examiner has three basic tasks: finding, preserving, and preparing evidence. Finding and isolating evidence to prove or disprove allegations is as difficult as preserving it. Investigators can plow through thousands of active files and fragments of deleted files to find just one that makes a case. Computer forensics has been described as looking

for one needle in a mountain of needles. Preserving computer evidence is important because data can be destroyed easily. The 1s and 0s that make up data can be hidden and vanish instantly with a push of a button. As a result, forensics examiners assume every computer has been rigged to destroy evidence, and they pro-ceed with care in handling computers and storage media .

Preparing evidence requires patience and thorough documentation so it can withstand judicial scrutiny. For example, a hacking incident at a Web music store was thrown out of court because examiners who prepared the case failed to follow rules of evidence that documented where evidence had come from and that it had not been changed .

Preserving computer evidence requires pre-incident planning and training of employees in incident discovery procedures. System administrators sometimes think they are helping a forensics examiner when they are actually destroying evi-dence. Managers should make sure that there's minimal disturbance of the com-puter, peripherals, and area surrounding the machine. If a computer is turned on, leave it on; if turned off, leave it off. Moreover, never run programs on a computer in question. For example, running Windows to examine files destroys evidence in the swap file. Finally, never let a suspect help open or turn on a machine .

Gathering computer evidence goes beyond normal data recovery. Unfortunately, there are no certified procedures for safe evidence gathering, nor is there a single ap-proach for every type of case. Examiners work in secure laboratories where they check for viruses in suspect machines and isolate data to avoid contamination.

Examiners will, for example, photograph equipment in place before removing it and label wires and sockets so computers and peripherals can be reassembled ex-actly in a laboratory. They transport computers, peripherals, and media carefully to avoid heat damage or jostling. They never touch original computer hard disks and floppies. They make exact bit-by-bit copies and they store the copies on a medium that cannot be altered, such as a CD-ROM. When suspects attempt to destroy media, such as cutting up a floppy disk, investigators reassemble the pieces to read the data from it. Nor do examiners trust a computer's internal clock or activity logs. The internal clock might be wrong, a suspect might have tampered with logs, or the mere act of turning on the computer might change a log irrevocably.

Before logs disappear, investigators are trained to capture the time a document was created, the last time it was opened, and the last time it was changed. They then calibrate or recalibrate evidence based on a time standard or work around log tam-pering, if possible.

Investigators always assume the worst. It is a rule in computer forensics that only the physical level of magnetic materials where the 1s and 0s of data are recorded is real, and everything else is untrustworthy. A suspect might have cor-rupted all of the software operating systems, applications, and communications in a computer, or the software itself might erase evidence while operating, so forensic examiners avoid it.

Examiners search at the bit level of 1s and 0s across a wide range of areas inside a computer, including email, temporary files in the Windows operating system and in databases, swap fields that hold data temporarily, logical file structures, slack and free space on the hard drive, software settings, script files that perform preset activ-ities, Web browser data caches, bookmarks, and history and session logs that record patterns of usage. They then correlate evidence to activities and sources.

Investigators have many tricks that help them get around the clever suspect. For example, they often do not attempt to decode encrypted files. Rather, they look for evidence in a computer that tells them what is in the encrypted file. Frequently, this evidence has been erased, but unencrypted traces remain to make a case. For data concealed within other files or buried inside the 1s and 0s of a picture, an investigator can tell the data is there even though it is inaccessible. Nearly identical files can be compared to see their minute differences.

When forensic examiners find computer evidence, they must present it in a logical, compelling, and persuasive manner that a jury will understand and a de-fense counsel cannot rebut. This requires step-by-step reconstructions of actions with documented dates and times, charts and graphs that explain what was done and how, testimony that explains simply and clearly what a suspect did or did not do, and exhibits that can withstand scrutiny.

Case presentation requires experience, which only can be gained through courtroom appearances. This is why lawyers and managers should retain computer forensics examiners who have a record of successful expert testimony on computer evidence. An experienced examiner knows the questions that opposing attorneys will ask and the ways to provide answers that withstand challenge. A skilled litigator can defeat an inexperienced examiner for failing to collect evidence in a proper manner and failing to show that evidence supports allegations. Not long ago, attorneys knew little about computers and how they operated, but today they do and they are increasingly skilled at challenging examiners' methods.

A Growing Service

With the growth of computers and networks comes the growth of crime commit-ted through or with computers and networks. It is a fast-growing field because computers and networks have moved to the heart of business and societal operations. However, it is not a service that most corporations will or should establish internally. Because investigations are so specialized, few organizations have the human or technical resources to gather and compile evidence that withstands court challenges. Large multinational corporations have or may develop the capability, but most organizations will purchase computer forensics as needed or keep a computer forensics firm on retainer. It's important that managers and lawyers remember that computer evidence is fragile and that the best way to handle an incident is to isolate it until examiners take over.

4.5 HIDDEN DATA AND HOW TO FIND IT

As if you didn't have enough to worry about, today's technology presents your business with as many problems as it does solutions. Computers that work miracles in your day-to-day operations often malfunction—and you lose valuable data. The email that makes communicating so simple, carries deadly viruses that infect your machines and spread, causing massive data losses throughout your network. Hack-ers, both inside and outside your company, can access your information, manipulate it, hide it, steal it, and cause huge losses of data.

In many cases, documents and files deleted from a computer can be found and recovered using the methods of computer forensics. When files or documents are deleted from a computer, the majority of the actual information is typically left be-hind. Although the user may think the deleted document has been eradicated, this is usually not the case.

Documents and files deleted or hidden even years ago may be recovered through a computer investigation. Deleted or hidden files are one of the prime tar-gets of the computer forensic technician searching for evidence.

What can you do about it right away? You should turn to computer forensic technicians or specialists (like Kessler International) for hard drive data recovery and other data recovery services. These technicians specialize in professional data recovery and will restore your data quickly—right when you need it. These teams of data recovery experts know how to retrieve your lost data from damaged and corrupt storage media including hard drives, back-up systems, temporary storage units, and more. They can also restore individual corrupt files back to their original condition.

SPYWARE AND ADWARE

Spyware is Internet jargon for advertising supported software (adware). It is a way for shareware authors to make money from a product, other than by selling it to the users. There are several large media companies that approach shareware authors to place banner ads in their products in exchange for a portion of the revenue from banner sales. This way, you don't have to pay for the software, and the developers are still getting paid. If you find the banners annoying, there is usually an option to remove them by paying the regular licensing fee.

Why Is It Called Spyware?

While this may be a great concept, the downside is that the advertising companies also install additional tracking software on your system, which is continuously *call-ing home*, using your Internet connection to report statistical data to the "mother-ship." While according to the privacy policies of the companies, there will be no sensitive or identifying data collected from your system and you shall remain anonymous, the fact still remains that you have a *live* server sitting on your PC that is sending information about you and your surfing habits to a remote location.

Are All Adware Products Spyware?

No, but the majority are. There are also products that display advertising but do not install any tracking mechanism on your system.

Is Spyware Illegal?

Even though the name may indicate so, spyware is not an illegal type of software in any way. However, there are certain issues that a privacy-oriented user may object to and therefore prefer not to use the product. This usually involves the tracking and sending of data and statistics via a server installed on the user's PC and the use of your Internet connection in the background.

What's the Hype About?

While legitimate adware companies will disclose the nature of data that is collected and transmitted in their privacy statement (linked from their database, there is al-most no way for the user to actually control what data is being sent). The technology is in theory capable of sending much more than just banner statistics and this is why many people feel uncomfortable with the idea.

On the Other Hand

Millions of people use advertising supported spyware products and could not care less about the privacy hype. In fact, some spyware programs are among the most popular downloads on the Internet.

Real Spyware

There are also many PC surveillance tools that allow a user to monitor all kinds of activity on a computer, ranging from keystroke capture, snapshots, email logging, chat logging, and just about everything else. These tools are often designed for parents, businesses, and similar environments but can be easily abused if they are in-stalled on your computer without your knowledge. Furthermore, these tools are perfectly legal in most places, but, just like an ordinary tape recorder, if they are abused, they can seriously violate your privacy.

4.6 ENCRYPTION METHODS AND VULNERABILITIES

The use of encryption provides a different kind of challenge for the forensic investigator. Here, data recovery is only half the story, with the task of decryption pro-viding a potentially greater obstacle to be overcome. Encryption, whether built into an application or provided by a separate software package, comes in different types and strengths.

Some of the most commonly used applications provide encryption protected by passwords that can be readily defeated by investigators with the right tools and the time to use them. Other types of encryption, readily available to the general public, can be configured and used to create encrypted data that goes beyond the ability of the professional investigator to decrypt it using software. Nevertheless, in these cases it may still be possible to decrypt data by widening the scope of the investigation to include intelligence sources beyond the computer under investiga-tion. For example, public key encryption can be used to create highly secure, encrypted data. To decrypt data encrypted in this fashion, a private key and passphrase is needed. The private key may be found on the suspect's machine or backed up to removable media. Similarly, the passphrase may be recorded some-where on the computer in case it is forgotten or may be written down somewhere and kept in a nearby location.

For example, three recent developments have added to the remarkable insecurity of the Net. One affects email. The other two affect network systems. First, a weakness has been discovered in the world's most popular encryption program that, in some circumstances, allows the encryption program to be completely by-passed. People using this program to encrypt email to protect its privacy and confidentiality may be thwarted despite their efforts. Second, hackers have recently discovered a cloaking program that allows them to blow past firewalls on servers and networks without being detected. Third, a flaw has been announced that affects networks around the globe regarding the file transfer protocol (FTP) used on the Internet. These three revelations taken together are seriously bad news for Internet privacy, confidentiality, and security.

The Fallacies of Encryption and Password Protection

How serious is the problem? Very. If a snoop can gain physical access to your com-puter or floppy disk where you store your secret key, he can modify it and wait for you to use it. When you do, he or she is secretly notified. From that point on, he has access to the rest of your encrypted personal information and you never know it. In effect, the snoop bypasses a user's *password* and bypasses the effects of encryption entirely. In this instance, the protection offered by encryption is illusory. Likewise, if a hacker can electronically break into your computer, and you have your secret key stored there, the security of your digital signature or your encrypted files is worthless.

Internet and Email Encryption and Security

For several years lawyers have been advised to use encryption programs to scramble sensitive email messages before sending them. The most popular encryption program is called PGP, or Pretty Good Privacy, invented by Phil Zimmerman a decade ago. PGP is a dual key, algorithm-based code system that makes encrypted data practically impossible to decipher. PGP is now owned by Network Associates, Inc. Of the 800 million people using the Internet, about 60 million use PGP to encrypt email.

In February 2001, Zimmerman went to work for Hushmail (an encrypted email system), aiming to make the use of PGP simpler and user friendly. His second goal was to work toward making PGP an international standard. To everyone's surprise, a month later, in March 2001, two engineers with a Czechoslovakian research group announced that they had found a serious flaw in the open PGP format.

The flaw is serious for two reasons. First, open PGP is the most widely used encryption system in the world. Until recently many systems that make e-commerce available by credit card on the Internet have been based on PGP. These products are still in use worldwide. Second, the theory behind PGP is essentially the same as that used in the Rivest, Shamir and Adleman (RSA) standard for digital signatures. The presumed *security* of this technique was what persuaded Congress to pass the Dig-ital Signatures Act, which is based on RSA standards.

Next, let's briefly look at how to protect data from being compromised. In other words, to protect data from being compromised, experts use computer forensics.

4.7 PROTECTING DATA FROM BEING COMPROMISED

In the past 25 years, since the introduction of the personal computer, a great change has taken place in the way people use computers. No longer are they an obscure rarity, but are ubiquitous, and the business without a computer is now an exception. They are used to assist with most tasks in the workplace. You communicate via email and chat, and even voice and video communication uses computers. You maintain financial records, schedule appointments, and store significant amounts of business records, all electronically.

It should come as no surprise that with this newfound productivity comes a class of individuals who exploit these benefits to commit crimes and civil wrongs. Almost any type of investigation and litigation today may rely on protecting evidence obtained from computer systems. Digital evidence can often make or break a case. This evidence may be used to establish that a crime has been committed or assert other points of fact in a court of law, such as identify suspects, defend the innocent, prosecute the guilty, and understand individuals' motives and intents.

As previously explained, computer forensics is the science whereby experts ex-tract data from computer media in such a way that it may be used in a court of law. In other words, computer forensics is used by experts to protect data from being compromised. This evidence may include such things as deleted emails or files and computer logs, spreadsheets, and accounting information.

It is not sufficient to merely have the technical skills to locate evidence on computer media. Computer forensics experts recover the evidence and maintain a strict chain of custody to ensure that the evidence is preserved in its original form. These experts' knowledge of what to look for and where to look is also important.

Now let's look at what has become a very popular undertaking: Internet tracing methods. How can you find out who is sending you email from a certain AOL or Hotmail account? Well, that's not what this next section is about is about. It is about how you can find out whether someone faked his or her email address and how you can find out from which account that mail really was sent.

INTERNET TRACING METHODS

If an email comes from a real, valid email account and you want to know who the person behind that email account is, then you most likely will need to serve the Internet provider who is hosting that email account a court-order. Another idea would be to take that email address and search for it on the Web and usenet. Who knows, he might have posted somewhere with his real name and address.

Sometimes people might send you information or hate mail from a fake address. This can be done quite easily by simply changing the Sender and Return-to fields to something different. You can do this, since these fields (your identity), are normally not checked by the mail server when you send mail, but only when you receive mail .

Every email has a so-called header. The header is the part in which the route the email has taken is being described. Since the header is rather ugly, it is normally hidden by the email program. Every email program can display them, though (look into the Options or Preferences menu).

The email text lines below are a typical, but not particularly sophisticated, ex-ample of faked email. Fortunately, most people are not more sophisticated than this. You should, however, be aware that there are much more sophisticated ways to fake email. A message sent to the newsgroup alt.security and archived (*http://catalog.com/mrm/security/trace-forgery.html*) on the Web explains one possible way to deal with some of these cases. But for now, back to the *easy cases*, as shown by the following email lines :

Received: from SpoolDir by IFKW-2 (Mercury 1.31); 13 May 98 15:51:47 GMT +01

Return-path: <kuno@seltsam.com>

Received: from bang.jmk.su.se by ifkw-2.ifkw.uni-muenchen.de (Mercury 1.31) with ESMTP;
13 May 98 15:51:44 GMT +01

Received: from [130.237.155.60] (Lilla_Red_10 [130.237.155.60]) by bang.jmk.su.se (8.7.6/8.6.6) with ESMTP id PAA17265 for <luege-ti@ifkw.uni-muenchen.de>; Wed, 13 May 1998 15:49:09 +0200 (MET DST) X-Sender: o-pabjen@130.237.155.254

Message-Id: <v03020902b17f551e91dd@[130.237.155.60]> Mime-Version: 1.0

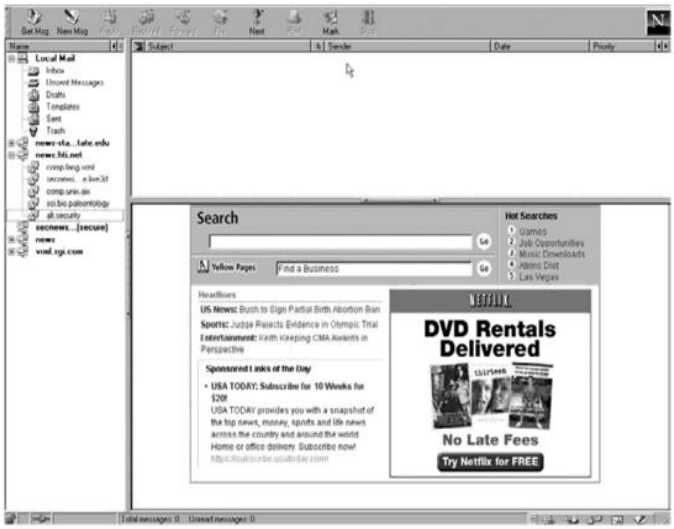

Figure 4.2 : A message sent to the newsgroup alt.security.

Content-Type: text/plain; charset=”us-ascii”
Date: Wed, 13 May 1998 15:49:06 +0200
To: luege-ti@ifkw.uni-muenchen.de
From: Kuno Seltsam <kuno@seltsam.com>
Subject: Important Information
X-PMFLAGS: 34078848 0

Now let’s go through the email line by line:

Date: Wed, 13 May 1998 15:49:06 +0200
To: luege-ti@ifkw.uni-muenchen.de
From: Kuno Seltsam <kuno@seltsam.com>
Subject: Important Information

The preceding lines should look quite familiar. They describe who claims to have sent the mail, to whom it was sent, and when. The following line is a number that your email program (in this case Pegasus Mail) might add to the mail to keep track of it on your hard disk:

X-PMFLAGS: 34078848 0

The following lines state that the message contains normal, plain text without any *fancy* letters like umlauts, etc.

Mime-Version: 1.0 Content-Type: text/plain;

charset=”us-ascii”

The following line contains a tracking number, which the originating host has assigned to the message. The Message-Id is unique for each message and in this case contains the IP number of the originating host. If you for some reason doubt that the message really came from someone at *seltsam.com*, you can now take this number and have it translated into something more meaningful. For this task, you can, for example, use TJPing (*http://www.topjimmysoftware.com/*), a small program that tracks IP packages online and resolves IP numbers:

Message-Id: <v03020902b17f551e91dd@[130.237.155.60]>

If you use TJPing, the real name of the originating computer is :

Starting lookup on 130.237.155.60 - May 14, 1998
22:01:25
Official Name: L-Red-10.jmk.su.se
IP address: 130.237.155.60

This is the originating computer from which the message was sent, not the mailserver. If the address was at a university, as in this case, this is not a great help, since there are many students using the same computers all day. The situation is very different within companies, though, since employees tend to have their own computers, which no one else uses. If the header doesn’t show any further information, you might use this information by calling the company’s system administration and ask, “Say, who’s sitting at Node 60?” Amazingly, often you will get a reply. It is comparatively easy to find out which company you are dealing with. Just cut off the first set of digits from the Official Name (L-Red_10.), add www and type it into your browser. You will see that *www.jmk.su.se* is the journalism department of the University of Stockholm.

The following line is solid gold. This tells you who was logged on to the mail-server when the message was sent. Not all email programs add this line, though. Eu-dora (*http://www.eudora.com/*) does, whereas Pegasus Mail doesn’t .

X-Sender: o-pabjen@130.237.155.254

So now you know that the user who sent us the mail is o-pabjen. The IP number is that of the mailserver used (checking with TJPing [*http://www. topjimmysoftware.com/*], *you* learn that it’s called bang.jmk.su.se). Now you could actually reply to the message by sending a mail to o-pabjen@130.237.155.254 or o-pabjen@bang.jmk.su.se.

Maybe you want to know his or her real name. In this case, you can try to *Fin-ger* the account. Finger is a command that reveals basic information about the ac-count holder. Due to the increased attention to privacy online, more and more servers have disabled it. It is always worth a try, though. Using WSfinger (*http://www.etoracing.com/wsfinger.htm*), you'll learn the following :

Login name: o-pabjen In real life: Pabst Jens global

So, now you have a name: Jens Pabst. *Global* could be part of the name or be some kind of code added by the system administration for internal purposes.

If you manage to obtain the information that's been accumulated so far, then you don't actually have to look any further. You have what you want. "Kuno Selt-sam <kuno@seltsam.com>" is really Jens Pabst <o-pabjen@bang.jmk.su.se>. But, let's go through the rest of the header anyway:

Received: from [130.237.155.60] (Lilla_Red_10 [130.237.155.60]) by bang.jmk.su.se (8.7.6/8.6.6) with ESMTP id PAA17265 for <luege-ti@ifkw.uni-muenchen.de>; Wed, 13 May 1998 15:49:09 +0200 (MET DST)

The preceding lines state which computer the mailserver has received the message from, when, and that the message is supposed to be sent to luege-ti@ifkw.uni-muenchen.de.

Similar to the last part of the header, the following lines tell you from where the recipient's mailserver (ifkw-2.ifkw.uni-muenchen.de) has received the message. You know that this must be the recipient's mailserver, since it is the last server that receives anything .

Received: from bang.jmk.su.se by ifkw-2.ifkw.uni-muenchen.de (Mercury 1.31) with ESMTP; 13 May 98 15:51:44 GMT +01

It follows the fake return path:

Return-path: <kuno@seltsam.com>

and an internal message from the mailserver about where and how it distributed the message within its system. You know that "SpoolDir" cannot be the recipient's mailserver, since it lacks an Internet address (something like server.somewhere.de).

Received: from SpoolDir by IFKW-2 (Mercury 1.31); 13 May 98 15:51:47 GMT +01

This next section is intended to familiarize the computer forensic investigator with various methodologies and tools available to perform a forensic examination of a Research In Motion (RIM) wireless (BlackBerry) device. The procedures and tools presented here are by no means all encompassing but are intended to elicit design of custom tools by those more programmatically inclined. The methods have been tested using an Exchange Edition RIM pager and an Exchange Edition RIM handheld.

4.8 **AVOIDING PITFALLS WITH FIREWALLS**

Consists of the pair of IP addresses that are talking to each other, as well a pair of port num-bers that identify the protocol or service. The destination port number of the first packet often indicates the type of service being connected to. When a firewall blocks a connection, it will save the destination port number to its logfile. This section de-scribes some of the meanings of these port numbers as well as avoiding some of the pitfalls. Port numbers are divided into three ranges:

- The well-known ports are those from 0 through 1023. These are tightly bound to services, and usually traffic on this port clearly indicates the protocol for that service. For example, port 80 virtually always indicates HTTP traffic.
- The registered ports are those from 1024 through 49151. These are loosely bound to services, which means that while there are numerous services "bound" to these ports, these ports are likewise used for many other purposes that have nothing to do with the official server.

- The dynamic and private ports are those from 49152 through 65535. In theory, no service should be assigned to these ports.

In reality, machines start assigning *dynamic* ports starting at 1024. However, there are exceptions: for example, Sun starts their RPC ports at 32768.

Suppose you are seeing attempts on the same set of ports from widely varying sources all over the Internet. Usually, this is due to a "decoy" scan, such as in "nmap." One of them is the attacker; the others are not.

Computer forensics and protocol analysis can be used to track down who this is. For example, if you ping each of the systems, you can match up the time to live (TTL) fields in those responses with the connection attempts. This will at least point a finger at a decoy scan. The TTLs should match; if not, then they are being spoofed. Newer versions of scanners now randomize the attacker's own TTL, making it harder to weed them out.

You can also attempt to go back further in your logs, looking for all the decoy addresses or people from the same subnets. You will often see that the attacker has actually connected to you recently, while the decoyed addresses haven't. A detailed discussion of firewall pitfalls.

Now let's briefly look at how both government and commercial organizations are implementing secure biometric personal identification (ID) systems to improve confidence in verifying the identity of individuals seeking access to physical or virtual locations for computer forensics purposes. In other words, a secure biometric personal ID system is designed to solve the fundamental problem of verifying that individuals are who they claim to be.

BIOMETRIC SECURITY SYSTEMS

The verification of individuals for computer forensics purposes is achieved using a recognized ID credential issued from a secure and effective identity confirmation process. A secure personal ID system design will include a complex set of decisions to select and put in place the appropriate policies and procedures, architecture, technology, and staff to deliver the desired level of security. A secure biometric ID system can provide individuals with trusted credentials for a wide range of applications from enabling access to facilities or secure networks, to proving an individual's rights to services, to conducting online transactions.

With the preceding in mind, biometric security systems for computer forensics purposes are defined as automated methods of identifying or authenticating the identity of a living person based on unique physiological or behavioral characteristics.

Biometric technologies, when used with a well-designed ID system, can provide the means to ensure that an individual presenting a secure ID credential has the absolute right to use that credential. Smart cards have the unique ability to store large amounts of biometric and other *data*, carry out their own on-card functions, and interact intelligently with a smart card reader. Secure ID systems that require the highest degree of security and privacy are increasingly implementing both smart card and biometric technology.

Finally, in an ID system that combines smart card and biometric technologies for computer forensics proposes to verify the identity of individuals, a "live" bio-metric image (scan of a fingerprint or hand geometry) is captured at the point of interaction and compared to a stored biometric image that was captured when the individual enrolled in the ID system. Smart cards provide the secure, convenient, and cost-effective ID technology that stores the enrolled biometric template and compares it to the live biometric template.

Unit V: Homeland Security Systems, Occurrence of Cyber Crime, Cyber Detectives, Fighting Cyber Crime with Risk Management Techniques, Computer Forensics Investigative Services, Forensic Process Improvement, Case Histories.

5.1 HOMELAND SECURITY SYSTEMS

Since 2000, terms such as "homeland security" and "homeland defense" have been widely used to describe America's response to the information warfare (IW) waged by terrorists . Let's look further into this last computer forensics system: homeland security.

Homeland Security Defined

The terms homeland security and homeland defense have received increased attention since the tragic events of September 11, 2001. While these terms are relatively new, the concepts behind them are not. Homeland security is defined as the deterrence, prevention, and preemption of and defense against aggression targeted at U.S. territory, sovereignty, population, and infrastructure as well as the management of the consequences of such aggression and other domestic emergencies. Homeland defense on the other hand is a subset of homeland security. It is defined as the deterrence, prevention, and preemption of and defense against direct attacks aimed at U.S. territory, population, and infrastructure. In other words, you might consider homeland security to encompass policies, actions, and structures designed to protect the rights and freedoms inherent in the U.S. Constitution and homeland defense a subset of homeland security with policies, activities, and actions designed to defend against extraterritorial threats, including preemptive operations. Nevertheless, the homeland security space is still being defined. A homeland security industry is still emerging.

Homeland Security Today

In November 2002, President Bush signed the Homeland Security Act of 2002, creating the Department of Homeland Security. The new department absorbs responsibilities from 22 agencies including the U.S. Coast Guard, Border Patrol, and Secret Service.

This is the most significant transformation of the U.S. government in over a half century. The creation of this cabinet-level agency is an important step in the president's national strategy for homeland security. The Department of Homeland Security has the following organizational structure:

- Border and transportation security
- Emergency preparedness and response
- Chemical, biological, radiological, and nuclear countermeasures
- Information analysis and infrastructure protection

Emergency mangers had been pleased with Bush's previous attention to emergency management. He was the first president to give the FEMA director an office in the West Wing. Now, emergency managers are concerned, as FEMA has been swallowed up in a new organization with a broader mission. Time will tell, but those who respond to and manage emergencies have much to do with the response to terrorist events.

The first line of homeland defense in any emergency is the "first responders"— local police, firefighters, and emergency medical professionals. Local first responders are the ones who will save lives and deal with the consequences of a terrorist attack. Emergency management and health care capabilities are a critical second tier to the first responders. While the U.S. is well prepared for "normal" emergencies, it does not currently possess adequate resources to respond to the full range of terrorist threats that are faced. Homeland security initiatives will likely focus on improving our capability to respond to a terrorist attack.

Emergency Managers and Homeland Security

Homeland security includes management of the consequences of terrorist acts and aggression and other domestic emergencies. This is the part of homeland security where first responders and emergency mangers play a vital role.

Emergency management is defined as a process to reduce loss of life and property and to protect assets from all types of hazards through a comprehensive, risk-based, emergency management program of mitigation, preparedness, response, and recovery. Emergency managers have been providing homeland security and homeland defense services for decades. During the Cold War this was called "civil defense" and the chief threat was a nuclear attack. Today, comprehensive emergency management, homeland security, and terrorism preparedness are included in an all-hazards comprehensive emergency management program (CEMP). Most emergency managers believe that homeland security should be included in a CEMP rather than developed as a separate program.

How Comprehensive Emergency Management Addresses Homeland Security

Finally, a CEMP is an overarching process that includes mitigation, preparedness, response, and recovery. A good program will address homeland security issues as well as continuity of operations, continuity of government, and related areas. Sound emergency management practices are required to mitigate the impact of day-to-day disruptions as well as managing response to and recovery from terror-ist attacks and other disasters.

5.2 OCCURRENCE OF CYBER CRIME

Cyber crime occurs when information technology is used to commit or conceal an offense. Computer crimes include:

- Financial fraud
- Sabotage of data or networks Theft of proprietary information
- System penetration from the outside and denial of service.
- Unauthorized access by insiders and employee misuse of Internet access privileges Viruses, which are the leading cause of unauthorized users gaining access to systems and networks through the Internet.

Cyber crimes can be categorized as either internal or external events. Typically, the largest threat to organizations has been employees and insiders, which is why computer crime is often referred to as an insider crime. For example, Ernst & Young's global research has found that 93% of all identified frauds were commit-ted by employees, almost 44% of which were committed by management.

Internal events are committed by those with a substantial link to the intended victim, for example, a bank employee who siphons electronic funds from a customer's account. Other examples include downloading or distributing offensive material, theft of intellectual property, internal system intrusions, fraud, and intentional or unintentional deletion or damage of data or systems.

However, as advances continue to be made in remote networks, the threat from external sources is on the rise. For example, in the 2003 CSI/FBI Computer Crime and Security Survey, 50% of respondents reported their internal systems as a frequent point of attack, while 59% reported Internet connections as the most frequent point of attack.

An external event is committed anonymously. A classic example was the Philippine-based 1999 "I Love You" email attack. Other types of external cyber crime include computer system intrusion, fraud, and reckless or indiscriminate de-liberate system crashes.

Internal events can generally be contained within the attacked organization, as it is easier to determine a motive and therefore simpler to identify the offender. However, when the person involved has used intimate knowledge of the information technology infrastructure, obtaining digital evidence of the offense can be difficult.

An external event is hard to predict, yet can often be traced using evidence pro-vided by, or available to, the organization under attack. Typically, the offender has no motive and is not even connected with the organization, making it fairly straightforward to prove unlawful access to data or systems.

5.3 CYBER DETECTIVES

Computer forensics, therefore, is a leading defense in the corporate world's armory against cyber crime. Forensic investigators detect the extent of a security breach, recover lost data, determine how an intruder got past security mechanisms, and, possibly, identify the culprit. Forensic experts need to be qualified in both investigative and technical fields and trained in countering cyber crime. They should also be knowledgeable in the law, particularly legal jurisdictions, court requirements, and the laws on admissible evidence and production.

In many cases, forensic investigations lead to calling in law enforcement agencies and building a case for potential prosecution, which could lead to a criminal trial. The alternative is pursuing civil remedies, for instance, pursuing breach of trust and loss of intellectual property rights.

Legal Issues

The most common legal difficulty faced by organizations seeking to redress cyber crime in the courts is having digitally based evidence accepted. Notwithstanding the technical expertise of information technology (IT) teams, most companies are illequipped to investigate cyber crime in a way that results in the collection of admissible evidence. For example, data collected as evidence must be shown to not be tampered with and be accounted for at every stage of its life from collection to presentation in court. In other words, it must meet the requirements of the jurisdiction's laws of evidence.

Another issue is the lag time between legislation and change and improvements in technology. As a result, law enforcement organizations and computer forensic experts are often forced to use archaic and nonspecific laws to fit unusual circumstances.

For example, to commit theft, a person must permanently deprive the victim of property. However, if a disgruntled employee copied an organization's database and sold it to a rival company, the organization is not permanently deprived of the data; therefore, technically, no offense of theft has been committed. In addition, it is unclear whether data fits into the legal definition of property. However, even in cases where there is a clearly defined crime, corporations are often hesitant to pursue a criminal conviction because of the time, cost, and reputation risk involved in reaching a legal outcome.

5.4 FIGHTING CYBER CRIME WITH RISK-MANAGEMENT TECHNIQUES

The rate of technological change, the spread of computer literacy, and the growth of e-commerce collaboration, such as alliances and marketplaces, make the challenge of restricting cyber crime damage daunting. With legislation lagging be-hind technology, businesses have had no choice but to absorb the responsibility for the security of their most valuable asset their information. Risks range from ex-pensive downtime, sales and productivity losses to corrupted data, damage to reputation and consumer confidence and loyalty, and hefty compensation payments or lawsuits for breaches of client information.

The best approach for organizations wanting to counter cyber crime is to apply risk-management techniques. The basic steps for minimizing cyber crime damage are creating well-communicated IT and staff policies, applying effective detection tools, ensuring procedures are in place to deal with incidents, and having a forensic response capability.

Effective IT and Staff Policies

Well-communicated and "plain English" IT policies educate staff about their rights and obligations in the workplace. The goal of these policies is to create a security solution that is owned by all staff, not only by those in the IT division. To be effective, IT policies should make plain what an individual employee can and cannot do on the organization's systems and the legal implications of misuse. It is also vital to make a continuing investment in policies, which must evolve and be supported by ongoing training initiatives.

Effective policies diminish the risk of internal attack, particularly unintentional attack. In addition, when attack does occur, these policies clearly define what constitutes a breach of security, making it easier to prosecute or seek compensation from the perpetrator.

Vendor Tools of the Trade

Although internal policies will not dissuade external cyber criminals, the right vendor tools will detect an external attack and alert the organization to the threat. These tools are programs that either analyze a computer system to detect anomalies, which may form the basis of an attack, or locate data that can be used as evidence of a crime or network intrusion.

Choosing the right cyber crime detection tools is essential for risk management in all organizations, but like most applications associated with an organization, the question is, what is the right tool? The right tools are those that deliver appropriate information that the forensic expert can interpret to achieve the best outcome. Ultimately, the evidence must withstand the rigors of legal proceedings. To deliver the information needed, software tools should be probing (without compromising the target of interrogation), concise, able to report findings fully, supported, and easy to use. Such tools will save forensic experts valuable time and allow them to concentrate on data interpretation.

The 2003 CSI/FBI Computer Crime and Security Survey shows a significant in-crease in companies using intrusion detection systems, from 58% in 2001 to 79% in 2003. Although some attacks will not be prevented, damage such as financial loss or negative publicity can be contained with early warning.

As with all of today's technology, detection tools date quickly as new threats emerge. Effective detection tools need to constantly evolve to counter these threats and must be engineered around best-practice risk management associated with vulnerabilities, system configurations, and viruses. Some online products and ser-vices currently on the market provide efficient, cost-effective solutions by accessing computer vulnerabilities specific to an organization's IT environment.

Effective Procedures

Even in an organization that has implemented the hardware, installed the software, produced the policies, and employed competent staff to run an effective IT environment, it is not possible to prevent an incident from occurring. However, the at-tack itself does not have the greatest impact on a company. How the business responds to that attack has the greatest impact on a company. Without the appropriate procedures in place to counter detected attacks, an organization is exposed to the risks of lost data, financial loss, network damage, and loss of reputation.

Although many different types of attacks may occur, the majority require the same basic steps of response. For example, the simple process of ensuring that the right people know about the incident when it happens enhances an organization's response, both in time and effective handling procedures.

Forensic Response Capability

When an incident occurs, an organization needs an appropriate forensic response in place. By appointing a forensic expert to manage the response to an incident, organizations ensure all avenues are canvassed, all evidence located and handled correctly, and all those involved treated impartially (see sidebar, "Computer Forensic Incident Response Procedures [CFIRP]").

In other words, deterrence is the appropriate forensic response and the fundamental element of a defensive strategy for the organization. However, for deterrence to be effective, potential antagonists must be convinced that they will be identified and punished swiftly and severely. This is the essence of the three key causal variables of general deterrence theory: certainty, severity, and celerity. Un-fortunately, while the methods for identifying perpetrators of crimes in the law enforcement context, and attackers in the military context, are well developed, similar capabilities do not currently exist for the networked cyber realm. Thus, while deterrence is recognized as a highly effective defensive strategy, its applicability to defense against attacks on our nation's information infrastructures is not clear, mainly because of our inability to link attackers with attacks.

A conceptual tool that can help visualize and understand the problem is to think of a thread, or sequence, of steps (with requisite technologies) necessary to effect a deterrent capability. As with the "weak link" and "picket fence" analogies, if any one of these steps is missing or ineffective, the ability to achieve the desired result is compromised.

Looking at this thread, you can see that current intrusion detection technology is focused primarily on the first element in the sequence above. Any response is generally limited to logging, reporting, and isolating or reconfiguring. What is missing is the ability to accurately identify and locate attackers and to develop the evidentiary support for military, legal, or other responses selected by decision makers. While defensive techniques are important, it's critical not to "stovepipe" in such a way that you can't effectively link with the offensive component of an over-all strategic cyber defense.

In addition to detecting the attacks, perhaps you should also develop a "forensic," or identification, capability to pass the necessary "targeting" information on to the offensive components of the response team, regardless of whether the response is through physical or cyber means. Such a capability is critical if your cyber defenses are to transcend a merely reactive posture to one in which both offensive and defensive techniques can be effectively applied in tandem. This is in line with the established principles of war, which suggest that an offensive (and therefore deter-rent) spirit must be inherent in the conduct of all defensive operations. Forensics response capabilities could help provide the bridge between the defensive and offensive elements of an overall cyber defense strategy. Accurate and timely forensic response techniques would also enable the effective use of the three elements of deterrence. Otherwise, attackers can act with impunity, feeling confident that they need not fear the consequences of their actions.

Forensics is a promising area of research that could help provide the identification and evidence necessary to support an offensive response against attacks on your information infrastructure, regardless of whether that response is executed through physical, information warfare (IW), or other means. Although forensic response techniques are highly developed for investigations in the physical realm and are being developed for application to computer crime, what is needed is an analogous capability for real-time, distributed, network-based forensic response analysis in the cyber realm. It would seem appropriate to incorporate the collection of forensic response data with the intrusion detection and response types of technologies currently being developed. Critical supporting technologies include those needed for correlation and fusion of evidence data, as well as automated damage assessment.

The importance of solid identification and evidence linking an attacker with an attack will be critical in the increasing complexity of the networked information environment. Cyber attacks against the U.S. and its allies may not have the obvious visual cues and physical impact typically associated with attacks in the physical realm. In these cases, the available courses of action will be heavily influenced by various political, legal, economic, and other factors. Depending on the situation, it may be necessary to have irrefutable proof of the source of the attack, the kind of proof typically developed through forensic response methods.

For example, one suggested concept is for a "cyberspace hot pursuit" capability to aid in the back-tracing of incidents to discover perpetrators. Use of such a capability implies the need for laws specifying authorization to conduct cyberspace pursuits and cooperative agreements with foreign governments and organizations. A second suggestion is for the development of a tamper-proof, aircraft-like "black box" recording device to ensure that when an incident occurs and is not detected in real time, the trail back to the perpetrator does not become lost.

Extending the aircraft analogy, the need for effective identification during cyberspace pursuits, and for coordinating offensive IW response actions through intermediary "friendly" networks, may necessitate a type of "network identification friend or foe (IFF)" capability, just as the introduction of fast-moving aircraft in the physical realm necessitated the need for secure IFF. Although the need for IFF has traditionally been a concern at the tactical level of warfare, the failure to effectively deal with such issues could certainly have strategic implications.

One issue of concern at the strategic level of IW is the distinction between the military and private sector information infrastructures. It is clearly not feasible to re-quire the private sector to secure its systems to the level required for military net-works. The approach suggested in this section may be applicable regardless of whether the networks attacked belong to the military. For example, in the physical realm today, if a civilian target is struck, the FBI and other federal agencies are called in to assist and investigate the incident, and when the identity of the attackers is de-termined, appropriate legal, political, or military actions are taken in response. From an organizational

perspective, efforts are underway to develop the necessary coordination structures, such as the National Infrastructure Protection Center, between the private and commercial sectors. From a technical perspective, major elements of the commercial infrastructure could participate in a national-level monitoring sys-tem, while private entities could maintain their own in-house capabilities with the ability to provide necessary data to national authorities following an incident just as would be the case with the FBI being called in to investigate a crime.

Another fundamental concern this approach may help address is the problem of malicious insiders. The security paradigm of enclaves separated by boundary controllers is most effective against attacks from the outside. Attacks initiated from within the enclave, possibly even by a trusted insider, have traditionally been much harder to defend against. Cyber forensics response techniques may provide the capability needed to deal with this problem, which simply cannot be addressed by tra-ditional security techniques based on privileges. These systems simply check whether a user is acting within the prescribed privileges while remaining in complete oblivion regarding the abuse of these privileges.

In other words, as previously discussed, a deterrence-based approach is an element of an overall cyber defense strategy. The need for timely and unequivocal identification of attackers is essential for such an approach to be effective. Unfortunately, the technical basis for such identification has not received much attention from the research and development community. In addition, there may be some complicating factors for the implementation of the type of identification and forensics response capability discussed here, such as the widespread move to encryption. However, until research and development resources are committed to investigation of the relevant issues, the extent of the challenge cannot be fully understood.

COMPUTER FORENSIC INCIDENT RESPONSE PROCEDURES (CFIRP)

Let's look at an incident that occurred at a well-known technical university that clearly shows the need to have an enforceable and workable CFIRP:

Picture this: it is 1 A.M. and email comes into the security mailing list from an outside source informing you that this site's server has been compromised, and from the logs two of the machines in your domain look to also have been compromised. The only people on the mailing list who are up and awake and reading their mail are the operations staff, but they know that sometimes in the wee hours, one of the more nocturnal network staff come in. They take a chance and call his office. To their delight, he is in his office, so they forward him the security email and consider their part of this incident finished.

The nocturnal network person reads the email, looks at the time, and decides to block those two hosts at the router from the Internet. He then sends an email to secu-rity stating that the hosts are blocked and considers his part in this incident finished.

The next morning the rest of the security team trickles in and reads the security mail along with about 500 other emails of various severities. The entire team assumes that the nocturnal network person notified the owner of the machines of the problem and that action has been taken. You all get on with other business, and of course the nocturnal network person, being nocturnal, is not around to correct your assumptions.

5.5 COMPUTER FORENSICS INVESTIGATIVE SERVICES

There are without doubt some very knowledgeable experts in the field of computer forensics investigations; however, there has been an increase in the number of people purporting to be experts or specialists who produce flawed opinions or take actions that are just plain wrong. The reasons for these errors are manifold but range from peer or management pressure, restricted timescales, and problems with soft-ware, to sheer lack of knowledge. Most investigations are basically the same in that they are either proving or disproving whether certain actions have taken place. The emphasis depends on whether the work is for the accuser or the accused.

In many companies, forensic computer examiners are *kings* because they have more knowledge of the subject than their peers. However, they are still subject to management pressures to produce results, and at times this can color their judgment. Time restrictions can cause them to take short cuts that invalidate the very evidence they are trying to gather,

and when they do not find the evidence that people are de-manding (even if it isn't there), they are subject to criticism and undue pressure.

Many of these *specialists* are well meaning, but they tend to work in isolation or as part of a hierarchical structure where they are the *computer expert*. The special-ists' management does not understand what they are doing (and probably don't want to admit it), and often they are faced with the question, Can't you just say this.....? It takes a very strong-minded person to resist this sort of pressure, and it is obvious that this has had an adverse effect in a number of cases.

This sort of pressure comes not only from within the organizations, but also from external sources. When you reply with: "I'm sorry it's just not there" or "No, the facts do not demonstrate that," you frequently end up with lengthy highpressure discussions with the client, which appear to be designed to make you doubt your own valid conclusions.

Working in isolation is a major problem; apart from talking to yourself (first sign of madness), many people have no one else to review their ideas and opinions. This is where having recourse to a team of investigators, software engineers, hardware engineers, and managers who understand (not always a good thing, depending on your point of view) any doubts or unusual facts, is valuable for fully discussing and investigating to ensure that the correct answer is found.

Computer Intrusion Detection Services

Installing technical safeguards to spot network intruders or detect denial-of-service attacks at e-commerce servers is prudent, but if your staff doesn't have the time or skills to install and monitor intrusion detection software, you might consider out-sourcing the job.

Intrusion detection is the latest security service to be offered on an outsourced basis, usually by the types of Internet service providers (ISPs) or specialized secu-rity firms that have been eager to manage your firewall and authentication. Al-though outsourcing security means divulging sensitive information about your network and corporate business practices, some companies say they have little choice but to get outside help, given the difficulty of hiring security experts [6].

For example, the Yankee Group reports that managed-security services (of which intrusion detection is the latest phenomenon) more than tripled, from $450 million in 2000 to $1.5 billion in 2003. By 2009, the market is expected to reach $7.4 billion, fueled by the trend toward outsourcing internal local area network (LAN) security to professional security firms as *virtual employees*.

Digital Evidence Collection

Perhaps one of the most crucial points of your case lies hidden in a computer. The digital evidence collection process not only allows you to locate that key evidence, but also maintains the integrity and reliability of that evidence. Timing during this digital evidence collection process is of the essence. Any delay or continued use of the suspect computer may overwrite data prior to the forensic analysis and result in destruction of critical evidence (see sidebar, "Evidence Capture"). The following are some helpful tips that you can follow to help preserve the data for future computer forensic examination:

- Do not turn on or attempt to examine the suspect computer. This could result in destruction of evidence.
- Identify all devices that may contain:
  - evidence: Workstation computers
  - Off-site computers (laptops, notebooks, home computers, senders and re-cipients of email, PDAs, etc.)
  - Removable storage devices (zips, Jaz, Orb, floppy diskettes, CDs, Sony Memory Sticks, Smart Media, Compact Flash, LS-120, optical disks, SyQuest, Bernouli, microdrives, pocketdrives, USB disks, firewire disks, PCMICA)

- Network storage devices (redundant array of independent [or inexpensive] disks [RAIDs], servers, storage area networks [SANs], network attached storage [NAS], spanned, remote network hard drives, back-up tapes, etc.)

- Quarantine all in-house computers:
  - Do not permit anyone to use the computers. Secure all removable media.
  - Turn off the computers.
  - Disconnect the computers from the network.
- Consider the need for court orders to preserve and secure the digital evidence on third party computers and storage media.

5.6 EVIDENCE CAPTURE

One of the fundamental principles of computer investigation is the need to follow established and tested procedures meticulously and methodically throughout the investigation. At no point of the investigation is this more critical than at the stage of initial evidence capture. Reproducibility of evidence is the key. Without the firm base of solid procedures that have been strictly applied, any subsequent antirepudiation attempts in court will be suspect and the case as a whole likely to be weakened.

Another frequent problem with capturing evidence is lack of experience not only lack of site experience but also inappropriate experience of the type of systems that might be encountered. One of the most difficult skills on-site is knowing when to call for help. It is essential that a sympathetic working environment is created such that peer pressure or fear of loss of status and respect does not override the need to call for help. Easier said than done perhaps, but no less essential for that reason.

Finally, sloppiness, time pressure, pressure applied on-site, fatigue, or carelessness have all been contributory factors in transforming solid computer evidence into a dubious collection of files. These avoidable issues come down to individual mental discipline, management control and policy, and selecting appropriate staff to carry out the work. They are issues for which there is no sympathy. This is bad work, plain and simple.

Ultimately, any time the collection of computer evidence is called into question, it is potentially damaging to everyone who is a computer forensic practitioner; it is ultimately in everyone's best interest to ensure that the highest standards are maintained.

Next, let's briefly look at drafting a comprehensive and effective computer forensics policy. This type of computer forensics service is used by countless organizations (banks, insurance companies, law firms, local governments, retailers, technology firms, educational institutions, charitable organizations, manufacturers, distributors, etc.).

Computer Policy

Often overlooked, detailed policies on the use of computers within an organization are an everincreasing necessity. Corporations and government agencies are racing to provide Internet access to their employees. With this access, a Pandora's box of problems is opened. Paramount is loss of productivity; workers can easily spend countless hours online entertaining and amusing themselves at their employer's expense. A hostile workplace environment can be created through pornography, potentially exposing the organization to civil liability.

Although protecting your organization from outside threats is clearly important, protecting the organization from internal threats is at least as important, if not more so. According to the 2003 Computer Crime and Security Survey con-ducted by the Computer Security Institute and the FBI, 67% of the respondents reported unauthorized access to information by persons inside the organization, compared to just 42% who reported intrusions by outsiders. A quarter reported theft of proprietary information, and 80% reported theft of laptop computers. Virus contamination was reported by 92%, and a staggering 99% reported systems abuse by insiders (pornography, pirated software, inappropriate email usage, etc.). According to Sextracker, an organization that tracks the online pornography trade, 82% of online pornography viewing occurs during the 9–5 work day.

Your computer forensics policy manual should therefore address all manners of computer-related policy needs. The content should be based on your corpora-tion's experience in employment-related investigations, computer crime investigations, civil litigation, and criminal prosecutions. Approximately half of the manual should consist of detailed discussions of each of the policy topic areas; the other half should be sample policies that can be readily customized for your organization. The discussions should include topics such as why policies are needed, potential liability, employee productivity considerations, and civil litigation. Safe-guarding critical and confidential information should be discussed in detail. The policies should directly address the problems that you would typically find in organizations of all sizes.

Now let's look at another computer forensics service: litigation support and insurance claims. As the risk increases, so will the interest in policies and the cost of premiums and litigation.

insure multimillion-dollar cyberspace policies. For carriers willing to sell such paper, the premiums have skyrocketed. Prior to September 11, 2001, the focus of information security was on critical infra-structure. After September 11, 2001, the focus has shifted to homeland defense and trying to understand whether financial institutions and other critical infrastructure such as telecommunications are vulnerable to cyber terrorism.

Insurance stalwarts such as Lloyd's of London, AIG, and Zurich now offer policies for everything from hacker intrusions to network downtime. The breadth of cyber insurance policies is growing, from simple hacker intrusion, disaster recovery, and virus infection to protection against hacker extortion, identity theft, and misappropriation of proprietary data.

While the market was already moving to provide policies to cover these risks, many executives viewed cyber insurance as a luxury that yielded few tangible benefits. Many risk managers buried their heads in the sand, believing they would never need anything like cyber insurance. There was a naiveté on the part of senior management. IT managers were not willing to admit they had to fix something of that magnitude, because they were afraid to go ask for the money.

The aftermath of the 9-11-01 attacks illustrates the interconnectedness of all systems: financial services, information and communications, transportation, electrical power, fire, and police. They all relate in profound ways we are only now be-ginning to understand. Businesses are starting to think about what type of recovery position they would be in if something similar to the World Trade Center attack happened to them.

While the cyber insurance market may grow in the wake of the 9-11-01 tragedy, carriers are tightening the terms and conditions of policies. Premiums are going up significantly and underwriters are hesitating to sign big policies. In the past, companies seeking a $25 million policy could find someone to cover them. Now it's much more difficult. Underwriters who didn't blink at $5 million or $10 million policies, would now rather insure $1 million policies. The marketplace is in transition, and there's undoubtedly a hardening of trading conditions for both traditional property and casualty insurance, as well as the emerging new e-commerce products.

Premiums on cyber insurance are an easy mark for price hikes because there's little historical data on which to set them. It's difficult to pinpoint the losses if data is corrupted, a network is hacked, or system uptime is disrupted. The fear of bad publicity keeps many companies mum on hacking incidents, which makes it more difficult to collect data for projecting future losses.

To develop robust cyber insurance, two major developments need to take place. First, sufficient actuarial data needs to be collected. Second, insurance carriers need to develop a better understanding of the IT systems in use and how they interact with other information and automated systems.

Industry analysts predict that underwriters will push any changes in cyber insurance offerings and the systems used by policyholders. The first indication of this trend came earlier in 2001, when an underwriting company tacked a 5 to 15% surcharge on cyber insurance premiums for users of Windows NT on Internet in-formation services (IIS) servers, citing their poor security track record, which makes them more expensive to insure. The underwriters are going to force the issue by saying, "Look, if you lose your whole business, if things like that happen, you can expect to pay a higher premium."

5.7 FORENSIC PROCESS IMPROVEMENT

The purpose of this section is to introduce the reader to a process that will enable a system administrator or information security analyst to determine the threat against their systems and networks. If you have ever wanted to know more about who might have attacked or probed your system than just the IP address that ap-peared in the var/log/messages of your machine, then this section may help you. Al-though it is rare, some of these simple techniques may help you identify the perpetrator of an attack on your system. Although most system administrators are rightly concerned with first securing their hosts and networks from attack, part of doing that job correctly demands that you understand the threat against those sys-tems and networks. The risk any system connected to the Net faces is a product of vulnerability and threat. The techniques covered in this section will help you de-termine possible actions and possible motivations of the attacker. If you can un-derstand your attacker, than you can better defend against and respond to attacks against your network. Of course, it is important to understand that hackers will loop through several systems during the attack phase.

So why bother researching the apparent source of an attack? What if your sys-tem is the first system of many that the hacker will use in his or her attack against other systems? Could you be held liable for damage done by the attacker to some-one else's systems? What if the attacker is operating from within a country that has no laws against hacking and can thus operate with impunity? Or what if the hacker is unskilled and has left clues behind that a skilled researcher could use to identify him or her? All of these reasons justify taking a small amount of time to research the apparent source of a serious attack or intrusion. Of course, all of these techniques should be used after you have secured your system and possibly consulted with law enforcement personnel. This should be done if the level and seriousness of the at-tack justify such an action. Next, let's review the tools that are used in the threat identification process.

The Tools

The tools discussed here outline a step-by-step process that will help you identify the attacking host and possible actors that may have used that host to attack your system. This section is not intended to be a tutorial for how to use each tool on its own. There are many sources of information that cover each tool by itself in more detail. Many of you are certainly familiar with or have used many of the tools dis-cussed here at one time or another. Keep in mind that here we are talking about the overall process of characterizing the threat from a domain. The first step in the threat identification process is simply to know who owns the IP used in the attack. For detailed switchology on the use of each tool, consult the main pages or other sources for each tool listed.

Dig –x /nslookup

The first step in the process is to reverse the offending IP address. The Dig -x ip command will perform a reverse lookup on an IP address from its domain name server. The "-x" option will ensure that you receive all records possible about your host from the Domain Name Service (DNS) table. This might include nameservers, email servers, and the host's resolved name. The "nslookup" command, Nslookup ip, will also perform a reverse lookup of the host IP address, but will only return the resolved name.

Whois

The next step in the process is to perform a "whois" lookup on the IP address to see who owns the offending IP or at least who it is registered to. This can be a tricky operation. Use the resolved name previously mentioned to try to determine what country or region the IP address might be based in and then be sure to use the proper who is gateway for that region of the world. The main gateways are ARIN (the American Registry), APNIC (the Asian Pacific Registry), and RIPE (the Euro-pean Registry). There are dozens of others, but most addresses should be registered in one of these. If your who is data does not match your resolved name, for exam-ple the resolved name http://www.cnn.com and the who is database ARIN indicates the registered owner is CNN network (a match), then you may have to do some more digging. Who is databases can contain outdated information. You may want to then research your IP with the country-specific who is database to determine the correct registered owner. A good collection of country-specific who is databases can be found at http://www.allwhois.com. For more information on conducting detailed who is queries check out http://www.sans.org.

Ping

Conduct the Ping ip command to determine if your attacking IP is currently online. Note that many administrators block ICMP traffic, so this is not conclusive evidence either way.

Traceroute

The next step in the process is to conduct a Traceroute ip to determine possible paths from your proxy site to the target system. Traceroute may help you in two ways. If your IP does not resolve possible paths from your proxy site to the target system, there may be a clue about its parentage. Look at the resolved host just before your target. This host's name may be the upstream provider for the attacking host and thus a point of contact or it may have the same domain as your attacking host, although that is not always true. Also, a traceroute might give you an important clue about the physical location of the attacking box. Carefully look at the path the packets traveled. Do they tell you what city they are in? Often they will. If you can determine what city the attack came from, you have just considerably nar-rowed down the possible pool of candidates of who the attacker might be.

Finger

Conduct a finger@ip command to determine who is currently logged onto the sys-tem that attacked you. Now, to be frank, this command will rarely work, because most administrators wisely turn this service off. However, it does not hurt to try. Keep in mind that many systems that are compromised and used as lily pads to at-tack other hosts are poorly configured (that is why they were compromised in the first place). They may also have the finger service running. If it is running, finger root@ip sees the last time root was logged on and, more important, from where root was logged on. You might be surprised to see root logged on from a third sys-tem in another country. Keep following the trail as long as your commands are not refused. You should be able to trace back hackers through several countries using this simple, often-overlooked technique. Look for strange login names and for users logged into the system remotely. This may indicate where the host was com-promised from and is the next clue to where to focus your research.

Anonymous Surfing

Surfing anonymously to the domain where your attacking IP is hosted is the next step in the threat identification process. You will know this domain name by look-ing at the resolved name of the host and the who is data. One technique that is use-ful is to use a search engine such as http://www.altavista.com with the specialized advanced search option of "+host:domain name and hack*." This query will return the Web links of possible hackers who operate from the domain name you queried. You can substitute warez or mp3 and the like to focus on terms of interest specific to warez or mp3 dealers. The number of Web pages returned by the query, as well as the details on those pages, gives you an indication of the level of threat to assess to a certain domain. For example, if you were investigating a host registered to demon.co.uk (Demon Internet), you would type "+host:demon.co.uk and hack*" in the Altavista query box. You may be surprised to see a return of some 55,000-plus hacking-related pages hosted on this domain. The Demon Internet seems to harbor many hackers and, as a domain, represents a viable threat to any organiza-tion. As a standard practice, you might want to block certain domains at your fire-wall if you are not already blocking ALL:ALL. Another possibility to widen the search is to use "+link:domain name" in the Altavista search. This will show all Web pages that have a link to the domain in question listed on their Web page. In other words, the ever-popular "here is list of my hacker friends and their c001 hacker sites" pages will appear via this search. You will also want to keep in mind the tar-get of the attack. What were the hackers going after? Can you tell? Conduct searches for the resources targeted and combine these terms with Boolean operators such as "and espionage." Check newswires or other competitive intelligence sources to determine, if possible, who might be going after your company's resources. A good site to use to conduct your searches anonymously is http://www.anonymizer.com.

USENET

The last step in the process of threat identification is to conduct a USENET traffic search on your domain. Sites such as *http://groups.google.com/* are excellent for this. Search on the attacking IP address in quotes to see if other people are reporting ac-tivity from this IP in any security newsgroups. Search on the domain name or hacker aliases that you might have collected from your anonymous surfing, or from the returns of your finger queries. You can expand the headers of the postings by clicking on "view original posting." This may show you the actual server that posted the message, even if the hacker attempted to spoof his or her mailing address in the visible header. This method can reveal the true location of your hacker. Clicking on "author profile" can also give you valuable information. Look at the newgroups your hacker posts to

and look at the number and sophistication of those postings. Pay attention to off-subject postings. A hacker will often let down his guard when talking about his favorite band or hobby, for example. You can also search sites such as *http://www.icq.com* if you have a hacker alias from a defaced Web page or your Altavista search narrowed by the domain "+hacker" criteria previously noted.

Putting It All Together

Once you have completed the process previously outlined and gathered all the in-formation from these tools, you should be able to make an educated guess about the threat level from the domain you are analyzing. With luck, you were able to collect information about the numbers and sophistication levels of the hackers who operate from the attacking domain, possible candidates for the attack (through finger or specialized Altavista searches), and what other CERTs may be seeing from that domain (via newsgroups or newswire searches). An excellent site to check for archived postings of recently seen attacks is both *http://www.sans.org* and *http://www.securityfocus.com*. *Ask* yourself, were there thousands of hacker pages hosted on the domain that you were investigating? Likewise, did you find thou-sands of postings concerning hacking on USENET? Did you run a search on your organization's name plus "hack*"? Were there postings from other administrators detailing attacks from this domain? Were the attacks they mentioned similar to yours or different? Now you might be able to determine if that FTP probe, for example, was just a random probe that targeted several other companies as well as yours or targeted your company specifically. Could you tell from the logs that the attacker was attempting to find a vulnerable FTP server to perhaps set up a warez or mp3 site? Being able to make an educated guess about the motivation of your hacker is important. Knowing whether your company has been singled out for an attack as opposed to being randomly selected will change the level of concern you have about assessing the threat. The process previously outlined can be used to narrow down possible candidates or characterize the threat level from responsible do-mains. As a byproduct, it will also provide you with all the necessary names, phone numbers, and points of contact that may be useful when it comes time to notify the pertinent parties involved.

Finally, let's look at what is probably the most important computer forensics service: training. It has now been expanded to support U.S. government and U.S. corporate needs, which became more of a priority after September 11, 2001. It places priority on computer incident responses and now covers computer forensic binary data searches for foreign language (non-Latin based) computer data (Farsi, Chinese, Japanese, etc.).

Training

As previously explained, computer forensics involves the preservation, identification, extraction, and documentation of computer evidence stored in the form of magnetically encoded information (data). Often the computer evidence was created transparently by the computer's operating system and without the knowledge of the computer operator. Such information may actually be hidden from view and, thus, special forensic software tools and techniques are required to preserve, identify, extract, and document the related computer evidence. It is this information that benefits law enforcement and military agencies in intelligence gathering and in the conduct of investigations.

Today computer forensics software tools and processing techniques have become important resources for use in internal investigations, legal electronic document discovery, computer security risk management, and computer incident responses. Computer forensic software tools and methods can be used to identify passwords, computer network logons, and other information that is transparently and automatically transferred from the computer's memory to floppy diskettes, Iomega Zip Disks, and computer hard disk drives. Such computer forensic software tools and methods can also be used to identify backdated files and to tie a floppy diskette to a specific computer. These techniques should be taught in your specialized training course.

Law enforcement and military agencies have been involved in processing computer evidence for years. Therefore, computer forensics training courses should be taught by certified instructors (see sidebar, "Computer Forensics Certified") who are experienced computer crime experts (retired federal law enforcement computer evidence trainers and members of law enforcement computer crime units).

Although most computer forensics training courses do not answer all possible questions regarding computer evidence and computer security, they should cover most of the common issues and expose the participant to new state-of-the-art com-puter forensics techniques and computer forensics tools. Training should consist of a Windows NT computer forensics course and a restricted-data-hiding course. An expert witness testimony on electronic evidence course should fill

in the gaps when the participant is ready for those advanced training courses. Training should not be focused on one specific computer forensics software tool or set of tools. This should not be a computer forensics paint by numbers training course. Quality computer forensic software tools should be provided with the training course, but it should be your company's mission to teach methodologies and the more technical aspects of computer evidence processing and computer incident responses.

The training course should be unique; the participants are expected to have a high degree of computer proficiency, know the difference between clusters and sec-tors, and have experience in the use of latest Microsoft Windows platforms. The course should not be an overview of computer forensics. It should be a technical hands-on training course that will tax your knowledge and computer skills. It should provide you with more information about computer security risks and evidence-processing information than can be found anywhere else.

Because the course should deal with computer security issues and computer risk management as well as computer evidence issues, it should be well suited for computer security specialists, computer incident response team members, and computer crime investigators. Most of your participants should be challenged by this course for it to be considered a success.

5.8 CASE HISTORIES

The following case study illustrates the organizational benefits of a planned forensic response.

Scenario One

An IT manager reviews a detection tool report that indicates a company employee is accessing restricted Internet sites and downloading objectionable material. After discovering the activity, the IT manager remotely accesses the employee's personal computer to obtain evidence. The employee is then dismissed, based on the evidence located and obtained.

Scenario Two

An IT manager reviews a detection tool report indicating a company employee is accessing restricted Internet sites and downloading objectionable material. After discovering this activity, the IT manager follows procedures, reporting his suspicions to the nominated computer incident response contact, in this case the chief information officer (CIO).

The CIO then invokes the company's incident response plan by contacting the incident response team, which includes computer forensics experts. This team isolates the *offending machine;* conducts a forensic examination of the computer sys-tem following methodologies known to be acceptable to criminal, civil, and arbitration courts or tribunals; and establishes where the material came from, how often, and who else knew about it. By following its effective policies and procedures, the organization (via the CIO) is in an excellent position to take immediate legal and decisive action based on all the available facts and evidence.

Which Scenario Works?

Only one of these scenarios illustrates a planned forensic response. In Scenario One, the evidence was obtained remotely. This fact alone may put the obtained evidence in doubt.

Any court of law would want to know whether there were policies and IT infra-structure for ensuring the IT staff member knew the correct PC was accessed. Other issues surround the need for evidence to prove that a particular employee's PC was responsible for downloading the objectionable material. Can it be proved that the objectionable material was viewed on a particular PC? Who else had access to that PC? It is likely that there is not adequate evidence in this scenario to answer these questions.

The IT manager detecting activity is only the first step in forming grounds for suspicion. If action is taken without proper policies, procedures, and processes in place, it is nothing more than an unplanned knee jerk reaction.

Unplanned reactions potentially expose an organization to risk. Clearly, any investigation must not only be thorough and methodical, but also staffs need procedures for reporting the activity, conducting the investigation, and appointing investigators.

Finally, in Scenario Two, the established policies let the organization clearly identify the incident and carry out appropriate immediate action. This places the organization in a comfortable position to resolve the situation, contain the potential damage, and effectively seek compensation or prosecution. The bottom line here is that without the appropriate procedures in place to counter detected attacks, an organization is exposed to the risks of lost data, financial loss, network damage, and loss of reputation.

5.9 COMPUTER FORENSICS CERTIFIED

According to a Gartner Group study, certification of INFOSEC computer-forensic-training professionals is becoming a common condition of employment. The research firm predicts that by 2009, INFOSEC certification will be required for 90% of CISOs (chief information security officers) and associated training staff positions and for 70% of day-to-day technical operations positions in Global 2004 companies. Security is the number one issue going forward in an online world, whether it's online voting or e-commerce.

THE DEMANDS OF SECURITY

It's bad enough when a certified IT employee doesn't possess claimed skills, but the skills gap is doubly worse in the security realm. What was once the near-exclusive purview of government agencies or companies involved in highly secret research is now a mainstream discipline for the highly connected enterprise.

This market didn't exist 13 years ago. The field has only matured in the past 6 years. Protecting a company's most cherished assets (not just IT systems, but especially the digitally stored proprietary information on those systems) demands knowledgeable personnel, something not always easy to assess. Anyone can hang out a shingle and say: "I'm an INFOSEC professional." Such people must be able to prove their credentials with INFOSEC certification.

Good security demands a more proactive approach than the other traditional functions of a system administrator. Security is the system administrator area that requires the most constant learning and relearning.

Information security infrastructure, like the proverbial chain, is only as strong as its weakest link. The breadth of skills and management ability required for strong information security puts unusual demands on organizations and professionals.

ANOTHER GAME

A certified information systems security professional (CISSP) isn't the only game in town. There's also Certified Internet Webmaster (CIW) professional certification, coming on strong.

Perhaps the best known security certification player is the System Administration, Networking, and Security (SANS) Institute, which sponsors the Global Information Assurance Certifications (GIAC). This is where the line in the security sand is drawn. The CISSP is a broad top-down certification, whereas the Level Two GIAC is a series of specialized and technical certifications.

GIAC responds directly to the skills question. GIAC requires that candidates demonstrate content mastery before they sit for the exam. In intrusion detection, for example, a candidate must accurately analyze ten real-world intrusion attempts before being allowed to take the exam. For firewalls, a candidate must design a perimeter that solves specific problems.

When comparing CISSPs to GIAC, the metaphor is an MBA (CISSP) versus a CPA (GIAC). You hire a CPA to do your accounting but not to do your strategic business planning. Research indicates that strategic business planning is what the industry desperately needs. The principal difference is in the target. An analogy suggested by an International Steering Committee (ISC) board member is that GIAC is for pilots and CISSP is for the managers who schedule the pilots.

SANS certification focuses on specific products. The product focus has limitations, because security professionals need to take into account the whole picture.

The short-term need is for the techie approach. Believe it or not, issues such as buffer overflows still form a large part of the action on security lists. In the long term, though, you need the big-picture view.

You cannot really say the technical issues are more important than management issues, but the technical issues are more solvable.

Indeed, whether approaching information security issues from a management or technical perspective, no one can escape political issues. Even if you had the best of the best techies on your payroll, you wouldn't be going anywhere unless the issues and policies around corporate standards, user awareness, remote and wireless access policies [8], acceptable authentication methods, and so forth have been decided. The critical success factors in most security jobs are being adept at the politics, possess-ing business skills and aptitude, good relationship management, and sales and negotiation skills, even in some lower-level jobs.

The product versus politics dilemma will eventually be moot with SANS' Security Essentials (LevelOne) certification. The basic GIAC certification now covers all the key knowledge sets covered by CISSP as well as additional, more current skills sets

Unit VI: The violation of privacy during information words. The individual exposed. Advanced computer Forensics systems and future directions advanced, encryption, hacking, advanced trackers, case studies.

ADVANCED ENCRYPTION: THE NEED TO CONCEAL

On German television several years ago, a stunned audience looked on as an un-suspecting Web surfer had his computer scanned while he was visiting a site. The site operators determined that a particular online banking program was installed on his computer, and they remotely modified a file in it so that the next time the user connected to his bank online, he also directed his bank (unbeknownst to him) to send a payment to the owners of that Web site.

The vulnerability of computer data affects everyone. Whenever a computer is connected to a network, be that a corporate intranet or the Internet, unless proper precautions are taken, the data residing in the machine can be accessed and other-wise modified by another knowledgeable user. Even computer data that the user may believe to be deleted or overwritten can be retrieved. Courts now routinely subpoena individuals' and companies' magnetic media as evidence; forensic experts can reconstruct data files that have been erased. In these cases, possession is not nine-tenths of the law. The best way to protect electronic data is to encrypt it.

The purpose of encryption is to render a document unreadable by all except those authorized to read it. The content of the original document, referred to by cryptographers as "plaintext," is scrambled using an algorithm and a variable, or key. The key is a randomly selected string of numbers; generally speaking, the longer the string, the stronger the security.

Provably unbreakable encryption has been around since the dawn of recorded history, and although computers have made encryption more accessible, they are certainly not a requirement. One precomputer method is the conceptually simple, yet very strong, encryption scheme known as the one-time pad, developed in 1926 by Gilbert S. Vernam (see sidebar, "Computer-Free Encryption").

COMPUTER-FREE ENCRYPTION

The durable encrypting scheme known as the one-time pad gets its name from the use of a key once and once only for just one message. It works like this:

Toni, the sender of a sensitive message wakes up one morning and starts shout-ing out two-digit numbers at random: 56, 34, 01, 92, 27, 11, and so on. These num-bers become the key. Toni then assigns a sequential number to each letter of the alphabet: A = 01, B = 02, C = 03, D = 04, E = 05, and so on.

Next, she encodes the plaintext word "hello," which, in accordance with the pre-ceding sequential numbering of the letters, corresponds to the sequence 08, 05, 12, 12, 15. She then does a simple modulo-10 addition with no carry, using the key she generated in the preceding. In other words,

```
  H  E  L L O
 08 05 12 12 15
+56 34 01 92 27
---------------
=54 39 13 04 32
```

This last sequence (54, 39, 13, 04, 32) is the cipher text, which gets sent to Wolf-gang, the intended recipient. Note that the same plaintext letters do not necessarily get encrypted into the same cipher text symbols (the letter L is both 13 and 04 in this case).

Wolfgang has an exact copy of the key (56, 34, and so on). To decode Toni's message, he does the reverse operation, again with no carry:

54 39 13 04 32
_ 56 34 01 92 27
_____
= 08 05 12 12 15
= H E L L O

Generating long keys by "shouting out" long strings of numbers can be impractical, so in modern applications of the one-time pad, computers are often assigned to create the keys. The result is not truly random: computers' pseudorandom number generators use only 16 (or, in some cases, 32) bits to store their values. The entire space of such values can be searched within a week or so. One remedy is to tweak the pseudorandom number generator by applying an external physical process to generate noise maybe a sufficiently amplified semiconductor junction of 1/f noise but that further requires removing the influence of predictable external influences, such as 50–60-Hz noise.

A self-evident shortcoming of the one-time pad is that the key is at least as long as the plaintext being encrypted. To escape cryptanalytic attacks involving statistical analyses, the key must be used only once. A more serious shortcoming is that the same key is used to both encrypt and decrypt. The sender and the recipient, there-fore, need a totally secure opportunity to exchange the key, which is hard to come by when the two are far apart.

An amusing feature of the one-time pad is that a fake key can be created that will "decode" the encrypted document into something quite innocent an excerpt from the Bible, say, or the Bill of Rights. Alternatively, a fake key could be designed to yield a plausible-looking, but still false, document, thereby fooling people into believing they have cracked the code.

Symmetric Encryption

Vernam's one-time pad is an example of symmetric encryption, in which the same key is used to both encode and decode a message. Many of the encryption schemes available today are also symmetric, most notably the Data Encryption Standard, or DES (see sidebar, "A Menu of Symmetric Encryption Algorithms")

A MENU OF SYMMETRIC ENCRYPTION ALGORITHMS

In symmetric encryption, the same key is used to encrypt and decrypt a message. Here are the most popular.

The Data Encryption Standard

DES was developed in the 1970s and is still used worldwide, although it has been replaced by the Advanced Encryption Standard (AES).

Triple Des

Encrypting the already DES-encrypted output with a different output with a different key provides no measurable security, but adding a third round of DES encryption yields a highly secure, albeit slower, algorithm. Most purportedly triple-DES implementations, however, use only two keys: key 1 for the first round of encryption, key 2 for the second round, and key 1 again for the third round.

The International Data Encryption Algorithm

The international data encryption algorithm (IDEA) uses a 128-bit key developed by ETH Zurich, in Switzerland. Its U.S. and European patents are held by Ascom Systec Ltd. of Bern, Switzerland, but noncommercial use is free. IDEA is viewed as a good algorithm for all except the best-funded attacks. It is used in Pretty Good Privacy (PGP) and Speak Freely (a program that allows an encrypted digitized voice to be sent over the Internet).

Blowfish

Blowfish is a 64-bit block code with key lengths of 32 to 448 bits. Developed in 1993 by Bruce Schneier of Counterpane Internet Security Inc., San Jose, California, it is used in over 100 products and is viewed as one of the best available algorithms.

Twofish

Twofish, also developed by Schneier, is reputedly very strong, and, as one of five candidates for AES, is now being extensively reviewed by cryptanalysts.

RC4

RC4 is a stream cipher of unknown security, designed by Ronald Rivest for RSA Security Inc., Bedford, Massachusetts. It adds the output of a pseudorandom num-ber generator bit by bit to the sequential bits of the digitized plaintext.

Developed in the 1970s, DES is still popular, especially in the banking industry. It is a block cipher, meaning that it encodes text in fixed-bit blocks using a key whose length is also fixed in length. The alternative, known as stream ciphers, encode the stream of data sequentially without segmenting it into blocks.

After nearly three decades of use, DES is headed for the garbage can. Currently, AES has already replaced DES in many organizations worldwide. In all likelihood, AES will become nearly as ubiquitous as its predecessor. Unlike DES, however, it will be competing with other algorithms algorithms that will not suffer from any suspicion that the U.S. government has a back door into the code.

For some encryption algorithms, a plaintext that is repetitive will result in a repetitive cipher text. This is clearly undesirable because the encrypted output betrays important information about the plaintext. One solution is to encrypt the cipher text block and add it bit by bit to the sequential bits of the previously encrypted plaintext.

Another problem with symmetric key encryption is that it requires that the sender and recipient of a message have a secure means for exchanging the encryption key. This is clearly difficult when the two parties are far apart, and the problem is compounded every time the keys are updated. Repeated use of the same key creates its own security weakness.

Public Key Encryption

An ingenious scheme that avoids many of the problems of symmetric encryption was proposed in 1976 by Stanford professor Martin Hellman and his graduate student Whitfield Diffie. Their public key encryption scheme, first described in IEEE Transactions on Information Theory, allows the recipient to verify that the sender is who he or she appears to be and that the message has not been tampered with.

The method works like this: Bob and Alice have a copy of openly available soft-ware that implements the public-key algorithm. Each directs his or her copy of the software to create a key, or rather, a pair of keys. A file encrypted with one key of a pair can only be decrypted with the other key of that same pair, one key cannot be mathematically inferred from the other key in the pair.

Bob makes known (by email, by posting to a Web site, or however else he chooses) one of the keys of his pair; this becomes his public key. Alice does the same. Each retains under tight control the other key in the pair, which is now his or her private key.

If Bob wants to encrypt a message that only Alice can read, he uses Alice's public key (which is available to anyone); that message can only be decoded by Alice's private key (Figure 6.1) . The reciprocal process (sending an encrypted message from Alice to Bob) is clear. In effect, Bob and Alice can now exchange encrypted

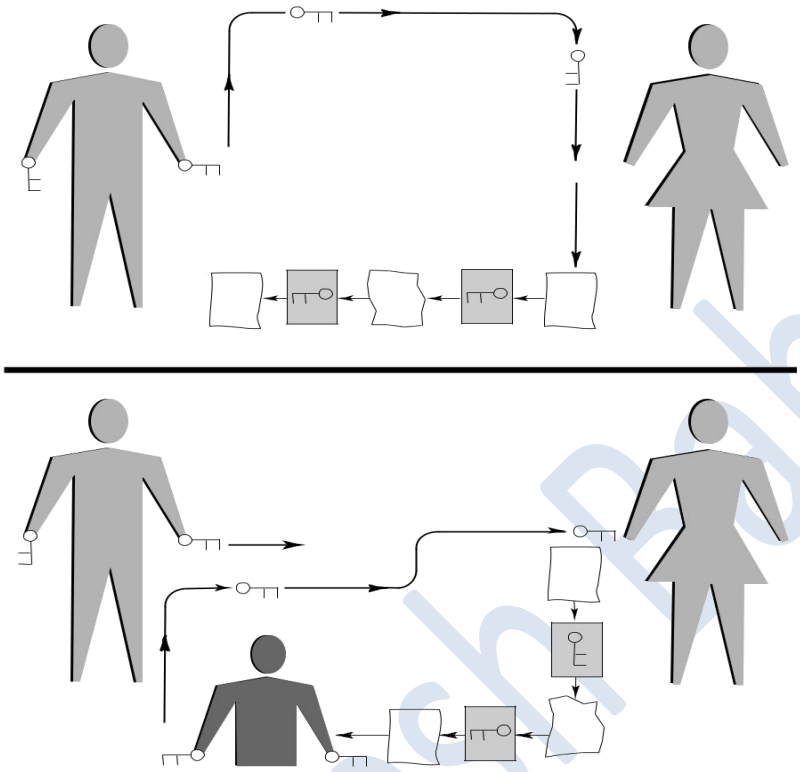

Fig 6.1

FIGURE 6.1 In public-key encryption [top], Alice encrypts a message using Bob's public key, and Bob decrypts it using his private key. This scheme allows encrypted files to be sent in the absence of a secure means to exchange keys, a major improvement over symmetric encryption. It's still possible, though, for Alice to receive a public key (or a conventional symmetric key) that ostensibly came from Bob, but that, in fact, belongs to a third party claiming to be Bob—the so-called man-in-the-middle attack (bottom).files in the absence of a secure means to exchange keys, a major advantage over symmetric encryption.

Sender authentication verifies that the sender is who he or she appears to be. Suppose Bob sends a message to the world after encrypting it with his private key. The world uses Bob's public key to decrypt that message, thereby validating that it could only have come from Bob.

Message authentication, the validation that the message received is an unaltered copy of the message sent, is also easy: Before encrypting an outgoing message, Bob performs a cryptographic hash function on it, which amounts to an elaborate version of a checksum. The hash function compresses the bits of the plaintext mes-sage into a fixed-size digest, or hash value, of 128 or more bits. It is extremely dif-ficult to alter the plaintext message without altering the hash value (Figure 6.2).

The widely used hash function MD5, developed by Rivest in 1991, hashes a file of arbitrary length into a 128-bit value. Another common hash function is SHA (Secure Hash Algorithm), published by the U.S. government in 1995, which hashes a file into a longer, 160-bit value.

Public-key encryption has been a part of every Web browser for the past few years. It is used, for example, when sending credit-card information to an online vendor or when sending email using the standard S/MIME protocol and a security certificate, which can either be obtained from online commercial vendors or created locally using special software.

One drawback of public key encryption is that it is more computationally intensive than symmetric encryption. To cut back on the computing, almost all

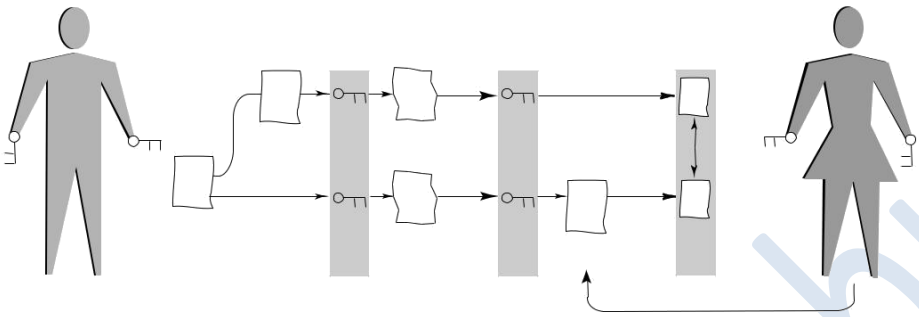

Fig 6.2

FIGURE 6.2 Public key encryption allows Alice to verify that a message from Bob actually came from him and that it is unaltered from the original. Here's how: Bob encrypts the hash value with his private key, encrypts the plaintext with Alice's (green) public key, and sends both to her. Alice then decodes the received cipher text using her own (orange) private key, decodes the hash value using Bob's public key, thereby confirming the sender's authenticity, and compares the decrypted hash value with one that she calculates locally on the just decrypted plaintext, thereby confirming the message's integrity. Implementations call on the symmetric approach to encrypt the plaintext and then use public key encryption to encode the local key. The differently encrypted plaintext and key are then both sent to the recipient.

In terms of resistance to brute force cryptanalysis (the exhaustive search of all possible decryption keys) a good 128-bit symmetric encryption algorithm is about as strong as a 2,304-bit public key algorithm. Realistically, though, the public key should be even longer than that, because the same public and private key pair is used to protect all messages to the same recipient. In other words, although a bro-ken symmetric key typically compromises only a single message, a broken public key pair compromises all messages to a given recipient. To be sure, cracking an encryption key is just one way to get at sensitive data (see sidebar, "Human and Hard-ware Frailties").

HUMAN AND HARDWARE FRAILTIES

The encryption of material to withstand a brute force attack still leaves many avenues open to invasion. Often, the real weaknesses in security lie in the human tendency to cut corners. It is all too tempting to use easy-to-remember passwords or keep unencrypted copies of sensitive documents on one's computer, intentionally or otherwise. Windows-based computers and many software products, in their quest to be user-friendly, often leave extensive electronic trails across the hard drive. These trails include not only copies of unencrypted files that the user deleted but also pass-words and keys typed.

Furthermore, unless each file is encrypted using a different key or a different encryption method, an attacker who can somehow read one encrypted file from or to a given person can probably also read many other encrypted files from or to that person.

Cryptanalysts have also been known to exploit the hardware on which the encryption algorithm is used. In 1995, the so-called timing attack became popular. It allowed someone with access to the hardware to draw useful inferences from the precise time it took to encrypt a document using a particular type of algorithm. Public key encryption algorithms such as RSA and DiffieHellman are open to such attacks. Other exploitable hardware phenomena include power consumption and RF radiation. It is also possible to assess the electronic paper trail left behind when the hardware is made to fail in the course of an encryption or decryption.

Most of today's commercial email programs, Web browsers, and other Internet applications include some encryption functions. Unfortunately, these schemes are often implemented as an afterthought by engineers who may be very competent in their respective fields but have minimal experience in cryptography. Just like a decent forgery, bad encryption can look like good encryption on the surface. In general, however, "proprietary," "secret," or "revolutionary" schemes that have not withstood the scrutiny of cryptanalysts over time are to be avoided.

One easy test is to attempt to decrypt a file with a different key from the one used for encryption. If the software proudly informs the user that this is the wrong key, that encryption method should be discarded. It means that the encryption key has been stored in some form along with the encrypted file. The cryptanalyst would merely have to keep trying different keys until the software identified the correct one. This is only one of many weaknesses. Given the preceding, the odds favor the person attacking an encrypted file, unless the person being attacked is very knowledgeable in the ways of information security.

Public key encryption is also a victim of the uncertainty besetting any crypto-graphic scheme when the two communicating parties lack a secure channel by which to confirm the other's identity (Figure 6.1). There is, as yet, no technical fix to this problem.

One of the most commonly used public key algorithms is the 24-year-old RSA, named for its creators, Ronald Rivest, Adi Shamir, and Leonard Adleman of the Massachusetts Institute of Technology, Cambridge. Its security derives from the difficulty of factoring large prime integers. At present, a key length of at least 1,024 bits is generally held secure enough. However, RSA may be somewhat vulnerable to "chosen plaintext attacks," attacks in which the cryptanalyst already possesses a plaintext file and the corresponding RSA-encrypted cipher text.

The DiffieHellman public key algorithm is used mostly for exchanging keys. Its security rests on the difficulty of computing discrete logarithms in a finite field generated by a large prime number, which is regarded as even harder than factoring large numbers into their prime-number components. The algorithm is generally viewed as secure if long enough keys and proper key generators are used.

By far the most popular public key encryption scheme is PGP (pretty good privacy). PGP was created in 1991 by a programmer and activist named Philip Zimmermann as a means of protecting email. After one of his colleagues posted PGP on the Internet, the Department of Justice launched an investigation of Zimmermann for possible violation of U.S. laws governing export of encryption products. The case against him was eventually dropped in 1996, after which Zimmermann started a company to market PGP. It has since become a mainstream commercial product, sold by Network Associates Inc., of Santa Clara, California, although freeware versions continue to be available from the Internet.

Crackdown on Cryptography

What happened to Zimmermann is just one small skirmish in the much wider campaign waged by governments worldwide against cryptography. At issue is whether, and to what extent, persons and organizations should have the ability to encrypt information that the state cannot decipher.

Private citizens have legitimate reasons to preserve confidentiality: to protect trade secrets, to prevent legal or medical records from falling into strangers' hands, and to voice dissenting political or religious opinions without retribution. The international group Human Rights Watch, for example, regularly encrypts eye-witness reports of serious abuse, gathered in parts of the globe where the victims may be subject to further reprisals.

From a government's perspective, however, encryption is a double-edged sword: it has honorable purposes, true, but it can also be used to conceal out-and-out criminality. In an effort to keep encryption from gaining ground, many countries have passed laws criminalizing its import, export, and use.

The proliferation of encryption has coincided with the explosive growth of the Internet. Nowadays, the man in the street can reach an instant global audience of mil-lions, bypassing the chain of command that rules almost any institution, be that the military, a religious group, or a corporation. In essence, the simultaneous spread of encryption and the Internet has amounted to a transfer of power to the individual.

This turn of events has been viewed differently by different states. An interesting case is the People's Republic of China. There, the outlawed religious sect Falun Dafa has used the Web to great effect to spread its ideology and recruit new members. Repeated attempts by the authorities to shut down the group have largely failed. Recently, the government began requiring any company doing business in China to disclose the types of Internet encryption software it uses, as well as the names of employees who use it. It further banned the sale of foreign-designed encryption products. Overseeing the regulations is a newly established body, the State Encryption Management Commission, which is believed to be staffed by China's secret police.

China's unwavering opposition to encryption suggests a more fundamental reason why a government (any government) would want to control the technology: to preserve the ability to exercise censorship. Even enlightened and democratic regimes have topics that are taboo, and when any and all information being ex-changed by private citizens can be monitored, it has a chilling effect on dissenting opinions. Conversely, when citizens can communicate freely and privately using encryption, censorship becomes unenforceable. Few sovereign states can accept this loss of control. It's like having two rude guests at one's dinner table who keep whispering in each other's ears.

Encryption, though, is good for business, and that factor is largely responsible for the gradual relaxing in the U.S. government's stance on encryption. Until 1996, strong encryption technology was listed as a munition, and until just recently, it fell under the same export restrictions as advanced weaponry. Under concerted pres-sure from the U.S. business community, which claimed that such controls were reducing sales and choking the growth of electronic commerce, the government came out with a revised policy recently that lifts many of the bureaucratic burdens from companies wanting to export encryption. Even so, every encryption product must still undergo a one-time review by the U.S. Commerce Department's Bureau of Ex-port Administration before it can be exported; sales to the so-called terrorist five (Cuba, Iran, North Korea, Sudan, and Syria) are still excluded. The new stipulation has some cynics wondering if only products with an identifiable weakness will receive an export license.

What's more (although encryption proponents have largely welcomed the relaxation of export rules), another concern has been raised: the same legislation would grant law enforcement new powers, such as the right to present a plaintext in court without disclosing how it was obtained from a suspect's encrypted files. Here the potential for abuse is obvious.

Other Legal Responses

The United States is not alone in backing away from strict encryption bans. What started as a global campaign to limit encryption has splintered into various approaches, with some governments now even encouraging encryption among their citizens as a precaution against snooping by other governments.

Generally speaking, laws pertaining to encryption are quite convoluted and rife with exceptions and qualifications. In Sweden, for instance, encryption importation and use are allowed, and so is its export, except to certain countries; authorities may search someone's premises for a decryption key but may not compel the person to assist in the investigation by, for example, handing over the key to the authorities.

The first international attempt to control encryption was made by the 17-coun-try Coordinating Committee for Multilateral Strategic Export Controls (COCOM), which came together in 1991 to restrict the export of items and data deemed "dan-gerous" if acquired by particular countries. COCOM members, with the notable exception of the United States, permitted the export of mass-market and public do-main cryptography and restricted export of strong encryption to select countries only. One such item was Global System for Mobile Communications (GSM) cellular telephony, which has two grades of encryption. Under COCOM, only the lower-grade version could be sent to the restricted countries.

In March 1994, COCOM was dissolved, to be replaced the following year by the multilateral Wassenaar Arrangement, which has now been joined by (at last count) 66 countries. Under the nonbinding agreement, countries agreed to restrict the ex-port of mass-market software with keys longer than 64 bits.

Do such encryption bans work? In a word, no. For one thing, the penalty for using encryption is likely to be far less than the damage caused by disclosing what-ever was deemed sensitive enough to warrant encryption. What's more, sophisticated techniques for hiding data, unencrypted or not, are now readily available and extremely hard to detect, so that prosecution of cryptography-ban violations is all but impossible. Who can prove that an innocuous-

sounding email message re-porting "The temperature in the garage was 86 degrees" really means "Meet me be-hind Joe's garage on August 6"? out of tens of millions of digitized images posted to a Usenet electronic bulletin board, who can detect the one image in particular, perhaps of an antique car, that contains a secret message intended for a specific per-son, who along with millions of unsuspecting others will download that image to his or her computer?

The very existence of the Internet has made it easy to circumvent bans. In most, though not all, countries, a sender can log onto any public computer connected to the Internet, such as those in public libraries or Internet cafés and send encryption soft-ware anonymously to a recipient, who can also retrieve it anonymously. A would-be user of encryption software can anonymously download it from any of the thousands of Internet servers that openly provide a large collection of programs of this kind.

It may make sense for a country to ban the exportation of something that it alone possesses and that could be used against it, but it makes no sense for a coun-try to ban the export of what other nations already produce locally. A 2004 survey by the Cyberspace Policy Institute of George Washington University, in Washing-ton, DC, identified 4,723 encryption products (hardware and software) developed in 82 countries.

The study, published before the latest relaxing of U.S. export laws, explains that on average, the quality of foreign and U.S. products is comparable and that in the face of continuing U.S. export controls on encryption products, technology and services, some U.S. companies have financed the creation and growth of foreign cryptographic firms. With the expertise offshore, the relatively stringent U.S. export controls for cryptographic products can be avoided since products can be shipped from countries with less stringent controls.

Nevertheless, in recent years, the war over encryption has moved beyond the mere control of the technology itself. Although encryption proponents may have won the first round, law enforcement and intelligence agencies have responded with a slew of powerful tools for getting at computerized data (encrypted or not). These efforts are in turn being met by ingenious new schemes for hiding and pro-tecting information, including one's identity.

ADVANCED HACKING

Today, as enterprise-wide networks reach the plant floor and zip data to the far side of the world in a twinkling, and, as the number of computers, personal digital assistants, telephones, and pagers communicating with the network increases, there is a corresponding increase in the opportunities for a critical blunder that would allow an attacker to enter your system. The consequences could be ruinous. Ac-cording to a recent survey by the Computer Security Institute, the cumulative loss of 520 companies that quantified their losses in 2004 reached $712 million, or about $1.4 million each. Roughly $595 million of that loss was theft of proprietary information—information your competitors want.

One of the best places for plant engineers to learn about network security (see sidebar, "Hack Yourself Before Somebody Else Does") with a peer in information technology (IT) is at the Computer Security Resource Center, a Web site (*http://csrc.nist.gov/*) established by the National Institute of Standards and Technology (NIST). There you'll find primers that explain security issues and technologies, news about current problems and security initiatives, and downloadable copies of the standards that govern electronic communication with Uncle Sam. More than ever, it's vital that plant engineers work effectively with IT to identify potential breaches, shore them up, and train everybody to be security conscious.

Nor is it only NIST that's getting into the act. The National Infrastructure Protection Center (http://www.nipc.gov/) was created by Congress to defend the nation's computer networks by serving as the national focal point for gathering information on threats to critical infrastructures. It is the principal means of facilitating and coordinating the federal government's response to an incident, mitigating attacks, investigating threats, and monitoring reconstitution efforts. The center issues updates about new viruses, Internet frauds, and disruption attempts almost daily. It is located in the FBI's Washington headquarters and maintains its own investigative staff.

Cyber security isn't an exclusively local matter, however. A complaint filed by the U.S. Attorney for the Southern District of New York provides an instructive exam-ple of the reach of today's e-thieves. The complaint alleged that Oleg Zezov and Igor Yarimaka, residents of Kazakhstan, penetrated the computers of Bloomberg.com,

in New York and demanded $200,000 from the company to tell how they had done it. Bloomberg agreed to pay, but only following a face-to-face meeting in London. There, accompanied by undercover London police officers, Bloomberg met with Zezov and Yarimaka. They repeated their demands, and police arrested them the next day. The United States is now seeking their extradition.

In view of the preceding incident, computer intrusions have more than tripled in the past two years. Who are the people trying to get their hands in your data, and why? Can you fight back by hacking yourself before somebody else does? This part of the chapter continues the theme of advanced computer forensics by providing answers to these questions.\

HACK YOURSELF BEFORE SOMEBODY ELSE DOES

How do you test your system to make sure it's as safe as possible? Can you recommend software, hardware, or services that can identify security issues before they become problems? What kind of procedures do you have in place to make sure that the latest patches are applied to Web servers?

The best way to retain your network security is to do frequent security audits, including trying to gain access using easily available hacking tools. In addition, you should ensure that you only run the services you need and only open the ports needed by your network.

Your gateway to the Internet should be a system without any important com-pany data or a hardware solution backed up by a firewall. You should also set up Windows Update notification for the server and have a backup server ready when you need to run the update.

Also, you should always check security bulletins and consider joining hacking mailing groups to find out what's happening on the other side of computer security. The main thing is to regularly test the security yourself; then you know what to find solutions for.

Are We a Hacker Nation?

Shadowy, computer-wise predators slip in undetected to steal data, deface Web sites, crash systems, or just look around. Why? Because hacking has become nothing in recent years if not a good career move. Yesterday's hackers are today's security gurus, with more corporations counting on them for protection.

One reason there are so many types of hackers these days is that hacking—at least as manifested in its simpler forms such as Web page defacement and denial-of-service (DoS) attacks (which overwhelm a site with data to prevent users from accessing it)—has never been easier.

Tools of the Trade

The Internet is filled with Web sites that offer tips and tools for the neophyte hacker. Kids, criminals, and terrorists are some of the people who avail themselves of this information—so more and more intruders are knocking at port doors. The barrier to entering the hacker world has become very low. If you have a political motivation against wheat farmers and you want to deface their Web page, you could just go online and learn how to do it.

Despite tighter Web security and stricter penalties for breaking into systems, hacking attacks have more than tripled in the past two years. The government's Computer Emergency Response Team reported about 39,000 cases of corporate hacking in the United States in 2002, more than 40,000 cases in 2003 and over 62,000 in 2004, and those are just recorded cases. To avoid negative publicity, most companies don't report attacks. The statistics cover network break-ins (which can give a hacker access to data files), Web site vandalism, DoS attacks, and data theft. The FBI estimates that businesses worldwide lost $5.9 trillion in 2004 from security breaches perpetrated from within the business.

The risks are personal and professional: hackers can steal passwords and bank account numbers from your home PC or grab trade secrets from your company network. Recently, criminal hackers broke into Microsoft's

corporate network and accessed source code for its software (see sidebar, "Future Threat: Advanced Mali-cious Code in Software").

Hacking also poses risks for national security—sophisticated terrorists or hos-tile governments could conceivably crash satellite systems, wage economic warfare by interfering with financial transfers, or even disrupt air traffic control.

FUTURE THREAT: ADVANCED MALICIOUS CODE IN SOFTWARE

Malicious code embedded in software is not new; users have always run the risk of downloading a virus or a trojan horse with shareware and games from the Net. The occasional intruder has even been found in shrink-wrapped products, but the hack into Microsoft's source code recently, raises worries that popular software may be the next target.

Although Microsoft indicates its code was not altered (the code was compared with previous backups) it's possible that a criminal hacker could get into a software manufacturer's code and insert a trojan horse. Unless software companies improve their security, you may find yourself the recipient of a gift horse in your next accounting package.

Good and Bad Hackers

Not all hackers have malicious intentions. Some hackers work for companies to se-cure their systems, and some contribute to security by notifying software vendors when they spot a vulnerability. Breaking things is easy. Building a solution is difficult, but arguably more fulfilling, but for every hacker who swaps his black hat for a white one, dozens of others continue to keep governments and companies on their toes.

Hacking will get worse. Bad software is being written faster than vulnerabilities are exposed. The trend is toward more features in applications, and the more features you have, the less security you get. Face it: hackers are not going to go away, so it's worthwhile to know who they are and why they do what they do.

Idle Hands

People see movies like *WarGames* and think hackers are going to start World War III. The truth is that computer hackers for the most part are smart, bored kids. Hackers usually start in their teens and stop by the time they're 30, but anyone can be a hacker—from the 16-year-old who defaces Web sites to the 36-year-old who sabotages a former employer's server. People in the underground indicate that not all hackers are true hackers.

By Any Other Name

It used to be that hacking had nothing to do with breaking the law or damaging systems. The first hackers, who emerged at MIT in the 1960s, were driven by a desire to master the intricacies of computing systems and to push technology beyond its known capabilities.

The hacker's ethic, an unwritten dictum governing the hacker world, indicates that a hacker should do no harm. A hacker should *pass through a network without a trace*. Somehow that message has gotten lost in the noise of Web defacements and data thefts.

Hacker purists get riled when anyone confuses them with crackers—intruders who damage or steal data, but although some hackers are quick to claim the moral high ground, the line between hacker and cracker is often blurred. Most hackers, for instance, don't believe it's criminal to break into systems and rifle around. The law, of course, thinks otherwise. Just because something is illegal doesn't mean it's wrong, but once you go in and destroy data or damage the system, that's where you stop being a hacker and you become a criminal.

T12, a 20-year-old who admits to some questionable hacking conduct, indicates he wouldn't normally damage a site, but if a phone company were to illegally switch his long-distance carrier and start billing his calls at $10 a minute, he wouldn't hesitate to take action. This is the kind of situation where one would feel free to just de-face their site and make it as public as possible.

Diablo, a teenager with the Romanian hacking group Pentaguard, indicates that a hacker should never abuse his or her powers, but if you penetrate a server and change the main page, nobody is hurt. The administrator gets embarrassed, and that's all.

Pentaguard has defaced more than 100 Web sites (most of them governmentand military-related) and Diablo indicates that he's careful: he never deletes or steals data and never crashes the system. This may be true, but the manager of one site Pentaguard defaced (owned by the Hawaii state legislature) indicated that his office had to pay $5000 for several new large-capacity hard drives (because the po-lice confiscated the hacked hard drives as evidence), and the site was down for a week until the drives arrived.

Signs of the Times

Hacking has definitely changed in the past 43 years. Talk to any hacker over 25, and he's likely to lament the passing of the good old days, when coding was an art form and learning how systems worked was an exercise in persistence. They say new hackers today are often younger and less skilled than their predecessors and more likely to focus on showy exploits than the noble pursuit of knowledge.

Many old hackers call the Internet generation of hackers *hollow bunnies*—such as gigantic chocolate Easter bunnies *filled with nothing but air*. Ten years ago, hack-ers respected information and machines and had to possess knowledge and skills to hack. Now novices use hacking programs without understanding them and are more likely to leave havoc in their wake.

Script kiddies receive the bulk of hacker disdain. These are the graffiti kids who download canned scripts (prewritten hacking programs) for DoS attacks or paint-by-number Web defacements. The risk here is that an unskilled hacker could re-lease wanton mayhem in your systems. The hacker might download a buggy hacking tool to your network that goes awry or execute a wrong command and in-advertently damage your machines. But script kiddies tend to disappear after a year. This is the generation of instant gratification, and if they can't get the hang of Back Orifice (a more advanced hacking program), they get bored and move on.

Bigger Threats

Script kiddies may get attention, but experts agree that the most dangerous hackers are the ones who don't make any noise: criminal hackers and cyber terrorists.

The truly dangerous people are hacking away in the background, drowned out by the noise and pomp that the script kiddies and DoS packet monkeys have been making.

Hacking has evolved into professional crime. Amateur hackers are falling into the minority, and now the fear is the criminal and the terrorist. These are people like the Russian cracker group that siphoned $20 million from Citibank in 1994 and the mafia boss in Amsterdam who had hackers access police files so he could keep ahead of the law.

In 1997, crime syndicates approached hackers to work for them. Now, with so many easy-to-use hacking tools on the Internet, criminals hardly need hackers to do their dirty work.

The cyber element that everyone fears most is one you've yet to see: foreign governments, terrorists, and domestic militia groups hacking for a political cause. The Department of Defense indicates its systems are probed about 583,000 times a year. It's difficult to tell if probes are coming from enemies seeking military data or from "ankle biters"—harmless hackers on a joyride. Regardless, authorities have to investigate every probe as a potential threat.

The likelihood of obtaining top secret information in this way is small, because classified data is generally stored on machines not connected to the Net. A more problematic assault would focus on utilities or satellite and phone systems. Ninety-two percent of U.S. military communications run through civilian phone networks. An attack on these systems could impede military communications.

For example, Navy officials recently reported that hackers broke into a Navy re-search facility in Washington, DC and stole two-thirds of its source code for satel-lite and missile guidance systems. The Navy indicates that the source code was an unclassified older version.

Thus, a large-scale cyber attack is imminent. Members of terrorist groups such as Hezbollah have been educated in Western universities and are capable of developing such attacks in the future—such as a digital 9-11 attack.

Why Hackers Hack

Aside from criminal and political motives, the reasons that hackers hack range from malice and revenge to simple boredom. Despite the image of hackers as dysfunctional loners, many are drawn to hacking by the sense of community it gives. Of course, a big part of hacking's attraction is the sense of power that comes from uncovering information you shouldn't possess. A hacker called Dead Addict once described the high that comes from discovering valuable information, followed by the low that comes from realizing you can't do anything with it.

For example, one hacker knows a little of that rush. He says that he once broke into a hazardous waste firm and found *pretty evil insider information* that no one was meant to see. Though he didn't act on the information, he did log it for possible use later—just in case he felt like being socially active.

Many hackers who begin as system voyeurs graduate to more serious activities. It's easy to be lured to the dark side when you get easy gratification messing around with individuals such as for example—AOL users. Most hackers are not old enough to drive a car or vote, but they can exert power over a network.

White Hats

There are a lot of the reasons why hackers' ability to hack into computers fade with age. Life fills their time and their ethics begin to change. The majority eventually find their interest waning. You only have three directions to go with hacking: you can keep doing the same old tricks, you can become a real criminal cracker, or you can use those skills wisely to build new software and create a more secure Internet.

Securing the Net is an interest many hackers develop (especially now that employers are hiring them for their skills). They lament that the public never hears about their positive acts, such as patching a hole on their way out of a site and let-ting the administrator know they fixed it. Most companies just focus on the fact that you hacked them and want to come after you with a lawsuit. It's made hackers reluctant to help them.

An even sorer point between hackers and vendors is the issue of releasing vulnerability exploits. These are findings about a security problem that hackers (and researchers) post on the Net. Vendors indicate hackers expose the holes for anyone to exploit and should instead report them to vendors first so they can fix them. The hacking community frowns on people who don't notify vendors, but when they do, vendors often ignore them. Most software companies won't do anything about a problem until you make it public. Then they have to fix it.

Vendors have a duty to develop secure software. Hackers, on the other hand, force vendors to admit their errors after they've hacked into their software. Manu-facturers are grossly negligent in selling software that doesn't stand up. What if they were producing cars that were this unsafe? The software they give us is not safe to drive in cyberspace.

Anything that's attached to the Internet is potentially hackable, and if you're using a Windows 2000, XP, or 2003 machine, nothing that is on that computer is secure. Better security is in everyone's best interest, and hackers should play a crucial role in this. The hacker kids who are going to Def Con today are the software architects of tomorrow. The same thing that makes them hackers makes them valuable to employers in the future.

All of this points to the fact that although hackers may be the Internet's greatest annoyance, their warnings are ignored about security at everyone's peril. The network that can't guard against a bored 18-year-old hacking in his or her spare time, can't hope to protect itself from a hostile government or tech-savvy terrorist.

ADVANCED TRACKER HACKERS

As the number of computer crimes spirals, the computer forensics experts' (a rare breed of security pros) skills are getting ever more precious. These are the data detectives who search for digital clues remaining on computers after malicious (or black-hat) hackers have done their dirty deeds. Cyber sleuths analyze email, Web site records, and hard drive data, looking for clues to the identity of criminals and crack-ers, much like gumshoes examine crime scenes for fingerprints and stray hairs.

It's not only the number of crimes that's fueling the need for these skills but also the increasing sophistication of criminals. The black-hat community is moving forward at a pace that outstrips the ability of the average system administrator or law enforcement agency. That means that both e-businesses and law enforcement agencies are paying plenty to find experts to sift through evidence left behind at dig-ital crime scenes.

In other words, security consultants and auditors are well-compensated for their knowledge—especially since the 9-11 attacks. In a recent survey of more than 11,000 IT managers, security consultants, on average, make $20,000 more per year than network administrators. Overall, salaries for all positions grew 33.9% to an average of $109,176 in 2004 (see Table 6.1).

TABLE 6.1 The Salary Protection Racket

Position	*Average salary*	*Increase from 2003*
Security consultants	$121,783	+23.7%
Security auditors	$110,722	+26.7%
Security administrators	$101,368	+29.9%
System administrators	$108,313	+44.3%
Network administrators	$104,047	+45.1%

The need for computer forensics is growing exponentially. The need is particularly acute at local, state, federal, and military law enforcement agencies that host computer forensics divisions, which are looking for individuals adept at solving hacking and intellectual property cases. An increasing number of corporations are using computer forensics to resolve internal matters such as fraud, violations of trade secrets, and inappropriate use of company computers.

The job is intense and tedious and requires nerves of steel. Most specialists have years of programming or computer-related experience, strong analytical skills, and the patience to invest days taking apart a computer in search of evidence. If things keep going the way they are, it probably won't hurt if these experts don't mind overtime.

Other professional attributes needed to catch a thief are strong computer science fundamentals, a broad understanding of security vulnerabilities, and strong system administration skills. Cyber sleuths use these skills to seek information to reconstruct how a system was hacked. The number and complexity of intrusions has increased at an alarming rate. Cyber sleuths have been forced to find ways to try to keep up with intruder tools as they have progressed in sophistication.

Experts gather this data and create an audit trail for criminal prosecutions. They search for information that may be encrypted or hidden, along with unallocated disk space. Most cunningly of all, they set traps using vulnerable computers to lure malicious hackers into giving away themselves and their techniques.

Computer forensics specialists must have strong analytic skills and excellent verbal and written communication skills. That's because they're required to document their findings in detail, and they often testify at criminal trials.

The demand is being answered by several educational facilities, including the University of Central Florida, in Orlando, which offers a graduate certificate degree in computer forensics. The International Association of Computer Investigative Specialists, based in Donahue, Iowa, offers certification for computer forensics examiners. Demand for such courses is so high that the association's fall classes are already full. Such courses are helpful for IT managers or individuals who lack computer programming experience but who want to make the leap into computer forensics.

Computer forensics specialists caution that IT managers interested in pursuing computer forensics as a career shouldn't expect that just by taking a few courses in the subject, they'll be able to track some of the world's slyest hackers (see sidebar, "The Costs of Tracking a Hacker"). The specialty is a tough discipline in a fast-moving industry that requires highly trained professionals dedicated to continued learning. That's because there's no way to stay ahead of the crooks. White-hat hackers at this point can only try to narrow the gap between themselves and the bad guys—and hope that the black-hat hackers don't get too fastidious when it comes to leaving behind digital footprints.

THE COSTS OF TRACKING A HACKER

It took the intruder less than a minute to break into the university's computer via the Internet, and he stayed less than a half an hour, yet finding out what he did in that time took researchers, on average, more than 34 hours each.

That inequity—highlighted during the Forensic Challenge, a contest of digital-sleuthing skills whose results were announced recently—underscores the costs of cleaning up after an intruder compromises a network. That damage done in a half an hour would take a company an estimated 34 hours of investigative time and cost about $3,000 if the investigation was handled internally and more than $33,000 if a consultant was called in—and those are conservative estimates.

Eventually, the members of a loose group of security experts known as the Hon-eynet Project, announced the winner of the Forensic Challenge. The contest pitted the reports of 13 amateur and professional cyber sleuths against one another.

Each digital detective used decompilers, data recovery programs, and other forensic tools to uncover as much information as possible. The entries consisted of a memo to fictional upper management, a security advisory, and an in-depth analy-sis of the evidence uncovered by the contestant's digital detective work. The winner of the contest, Thomas Roessler, a student in mathematics at the University of Bonn in Germany, has dabbled in, but not done digital forensics work in the past. Roessler indicated that it's always amazing how much information you can get out of a sys-tem by using rather basic tools. You always miss something.

The contest was made more interesting by the fact that the attack was a real one, captured by one of the several "honeypots" (vulnerable computers con-nected to the Net and surreptitiously watched) run by the Honeynet Project. In fact, the detectives produced several leads to the identity of the culprit. However, the person responsible would not be prosecuted. Such on-line vandals are extremely common.

The perpetrator represents a very large and common percentage of the black-hat community. It's a threat that everyone faces. Nevertheless, only about 70 to 80% of the so-called black-hat hackers (those who break into computers illegally) have com-parable skills to the attacker who breached the computer.

The contest also helped illuminate why securing a computer is more cost-effective than hiring consultants to come in and do the detective work afterward. It is a fairly extensive process to take what amounts to a bunch of garbage and build a comprehensive picture of what happened. The costs of such investigations can easily amount to $63,000 per computer.

Companies need to understand the difficulty, and costs, involved. Companies also tend to balk at agreeing to that kind of expense when there is no guaranteed pay-off. Maybe the contest opened the eyes of corporate executives, who all too often want a quick fix.

If you just reinstall the system, do you know if you have plugged the hole that allowed the attacker to get in? Most of the time, such quick fixes just mean the attacker gets another shot at the system. Some computers at the University of Washington have been compromised five times. Multiple intrusions are occurring all over the place.

The Honey net Project plans to do another contest soon, but it's a question of time. The next project would also focus on either a Solaris or Windows NT/2000, XP, or 2003 computer.

Anonymity in Retrieving System Logs

The 9-11 terrorist attacks have had numerous effects on national security. One of these is legislation that increases the ability of federal agencies to intercept Internet traffic. Another side effect was the loss of the well-known Web anonymity service hosted by ZeroKnowledge, which turned out not to be related to any of the national security activities during 9-11.

Web anonymizers allow people to visit Web sites without disclosing their iden-tities to the owner of the Web site, or even a local administrator who can log the URLs that a user visits. These tools work just as well for a terrorist who wishes to use the Web with anonymity, although using Internet access in a Web café, which one of the plane hijackers did, works well, too.

Anonymity has its place in a free society, and personal rights and freedoms shouldn't be collateral victims of terrorist attacks. Interestingly, government agencies may also be important users of anonymizers. This section explains how anonymization works on the Internet and why this is important in the face of in-creasing privacy concerns.

Source Addresses

A source address is an IP address embedded in the header of an IP packet. When the packet is received, the source address becomes the destination address in the reply packet. If you spoof your source address, reply packets wind up going to the address that you've spoofed, and you don't see the results. Worse, spoofing your source address is a lousy technique for anonymity, as most application protocols require a completed transmission control protocol (TCP) connection before exchanging any information.

Something similar to source address spoofing occurs whenever a firewall is between you and the destination network. Most firewalls translate internal addresses into external addresses, most commonly through network address translation (NAT). Another way to rewrite the source address is to connect to a proxy and ask it to connect to the server you want to visit. This capability is built into Web browsers, which permits you to specify the IP address and the proxy you wish to use. If you've configured your Web browser to use a proxy, the Web server sees the address as the source. The proxy relays for you transparently.

Of course, whoever maintains the proxy or firewall has logs of your activity. And the owner of the Web server still has information about you—for example, the type of Web browser you're using, the source IP address, the URL requested, any refer-ring page, as well as the source operating system, and sometimes the type of PC.

The Web site's operator can go farther still: a Web designer can include Javascript to collect more information about your browser and operating system, attaching that to any form data that you return. This information may include your system's real source IP address, which is accessible to Javascript programs.

Routing to the Rescue

Surprisingly, the U.S. Navy has researched network anonymity. This research formed the basis for the Freedom Network and may show up in other systems for anonymity as well.

Suppose that you proxy your Web requests through a third party who promises to keep its logs a secret. You connect to this server via secure sockets layer (SSL) so that anyone sniffing the connection can only see that you're visiting an anonymizer, and not your final destination site, which is encrypted. Sounds like a reasonable so-lution, but it hasn't worked in the past.

In the early 1990s, a site in Finland, *anon.penet.fi*, provided an anonymous re-mailer. Anonymous emailers strip away revealing information from email headers before resending it to your intended destination. That works well as long as the software manages to remove all the headers and you don't include revealing information in the email you send (for example, including an automatic signature file at the end of your email that, consequently, identifies you).

Penet also supported using aliases, so that the person receiving your email could reply to you without learning your identity. Therefore, Penet had to keep track of the mapping between your anonymous email address and your real one. Penet worked well until authorities stepped in and demanded that Johann Helsingius, Penet's operator, disclose the mapping of a particular email address be-cause it involved information copyrighted by the Church of Scientology.

If the proxy doesn't even know your real source address, how can it successfully relay for you? There have been several approaches to this problem, and one of the most recent (as previously discussed) is Onion Routing.

In Onion Routing, instead of having a single proxy for relaying, there's a net-work of proxies. Each of these proxies runs the same software, which not only re-lays your packets but also encrypts them. The first Onion Router chooses a route for your connection, then encrypts your data several times, each time using the public key for one of the routers in its network of routers.

This is where the "onion" comes in. Each layer of encryption resembles the skin of an onion: the Onion Router you've connected to first encrypts your data using the key of the last router in its list of routers—this makes up the innermost layer of the onion. Once this layer of encryption is removed, the packet is sent to its real destination. Then, the first Onion Router adds another layer of encryption. This layer includes the address of the last router in the list, and gets encrypted with the second to last router's key. The next layer gets added, with the address of the sec-ond to last router's address, but using the third to last router's key, and so on. There should be at least six routers to ensure confidentiality.

Onion Routing is even more effective if you run one of the routers. Your Onion Router must also be a full participant in the network, so that other Onion Routers can use it. Otherwise, packets coming from your Onion Router will only contain packets from your network, and can reveal your approximate source, even with the content still encrypted.

Onion Routers present another potential problem. An aggressive attacker could monitor the network traffic of every participating Onion Router. This at-tacker (or snoop) can then track traffic patterns. For example, you send off a re-quest to *http://www.fbi.gov* via your Onion Router. The snoop sees traffic leaving your Onion Router, bound for another Onion Router, with a certain packet size. The next router sends off a slightly smaller packet and so on, until the final router sends the plaintext packet directly to the real destination. Then the snoop can de-duce that this packet came from your network, based on the sizes and the timing of the packets between routers.

Onion Routing defeats this by delaying packets slightly, as well as batching data from several packets. Thus, a snoop cannot make simple deductions about the size and timing of packets. The end user does experience greater latency (delay), but this is the price for greater security.

Onion Routing is only one approach to the problem of network anonymity. AT&T Research (*http://www.research.att.com*) tried a different approach called Crowds. The concept behind Crowds is that "anonymity loves company," so the more participants the better. Each Crowd proxy is called a "jondo" (think "John Doe"). Unlike Onion Routing, which relies on layers of encryption, jondos employ secret key encryption with one key per route. This speeds up processing by reduc-ing the amount of time required to handle encryption. As with Onion Routing, state information is required so that the entry and exit points of a route know where to send packets. This information is discarded at the end of each connection but could be used to track users.

The Freedom Network used an approach similar to Onion Routing. You could either add a plug-in to Internet Explorer or patch your Linux kernel so that your sys-tem actually becomes an entry point in the network, with sites other than the one run by Zero Knowledge participating as routers. The Freedom Network claims that it decided in spring 2001 to discontinue its service because it wasn't paying for itself.

As of this writing, the Anonymizer (*http://www.anonymizer.com*) is still up and running but functions as a proxy; it also strips identifying information from your requests. Although you can use this service for free, your request will be delayed so that you can read ads encouraging you to pay for the service.

You can also acquire software that acts as a local proxy for Web requests. This software removes the USER-AGENT line and strips away cookies, which can also be used to track your use of a Web site.

Who Needs It?

The Onion Routing project closed down in January 2000, after processing over 30 million requests. Its home page contains an interesting disclaimer, essentially saying that anyone using the Navy's network should expect their traffic to be monitored— a very chilling statement when one considers the alleged intent of Onion Routing.

Still, government agencies form one of the largest groups of anonymizer users. Anonymizers allow law enforcement to visit Web sites without giving away their identity, or military analysts to collect data without revealing their areas of interest. Such uses of anonymizers are legitimate and actually of value to national security. If only the military and law enforcement used a particular anonymizer, then any visits from that anonymizer would immediately be of interest to someone worried about being investigated.

Anonymizers also have a place for nongovernmental users. While an anonymizer has the potential for misuse—for example, by hiding the identity of visitors to a pornographic site with illegal content—anonymizers have historically had more important and legitimate uses. For example, someone with AIDS could feel free to search the Web without revealing his or her identity. A person on the verge of committing suicide could ask for help, while remaining anonymous, which was one of the actual uses of the original Penet remailer. One can only hope that the rush to embrace national security in the United States doesn't have additional ca-sualties—especially ones that actually enhance national security.

Denial of Service

Pity the poor intrusion detection system (IDS)—it has the reputation of an irritat-ing snitch and the track record to prove it. Perhaps no other security device has done its job so well and then been reviled so roundly for doing it. Designed to sniff out and warn system administrators when hackers are trying to exploit network vulnerabilities or launch denial-of-service (DoS) attacks, the original IDSs did their job all too well. That was both bad and good news.

True to vendor promises, first-generation IDSs generated information-traffic patterns on network segments, aberrations in host log files, and so forth, which could indicate whether their systems had been hit with any of the attacks hackers use to break into critical network resources. This required placing IDSs at key locations on the network, such as at firewalls, switches, routers, Web servers, databases, and other back-end devices further into the enterprise—a straightforward process.

Those IDSs were also overly chatty boxes, renowned for generating mountains of data on traffic passing through networks and on host systems. They cried "wolf " too often, reporting false alarms by the droves. Consequently, many systems ad-ministrators, overwhelmed by tons of information they couldn't digest or didn't understand, simply dumbed them down or shut them off entirely.

The IDS products on the market are now bigger, better, and faster and offer much more to those charged with protecting network resources. Vendors have, for instance, developed new intrusion detection methods that go beyond the pattern, or signature-matching, technology that plagued the earlier products with all those false alarms. They have also increased the performance of their devices, which can now keep up with 100 Mbit/sec networks. Vendors are also shipping appliance-like IDSs, which simplify their deployment and management, and they've begun delivering products that combine the best of the two principal types of IDSs into a single offering.

Just as importantly, the number of attacks on networking systems is growing. It's a jungle out there, and network managers need to keep the predators at bay with a variety of security devices, including the IDS. For example, the nonprofit CERT Coordination Center received reports on 44,304 security incidents in 2004 (the most recent year for which its incident totals are available). That's comparable to the 33,879 it received in 2003, and the 22,768 incidents logged for 2002.

The most virulent threat to emerge from the hacker jungle, though, is clearly DoS and distributed DoS (DDoS) attacks, the number and variety of which have in-creased dramatically according to security organizations. Hackers target DoS at-tacks at devices and networks with Internet exposure, especially e-commerce sites [4], according to the National Infrastructure Protection Center (NIPC). The goal of such attacks is to incapacitate a device or network with bandwidth (devouring traf-fic so that external users can't access those resources)—this without hacking pass-word files or stealing sensitive data.

In March 2004, NIPC continued investigating a series of organized hacker ac-tivities that specifically targeted e-commerce and online banking sites. NIPC iden-tified 500 victims in 33 U.S. states who were attacked by organized groups in Eastern Europe (particularly Russia and the Ukraine), which took advantage of vulnerabilities in servers running an unpatched version of Microsoft's Windows NT operating system. Once the Eastern European hackers gained access, they downloaded a variety of proprietary data—mostly customer databases and credit-card information. In this case, the intruders didn't use the information maliciously, per se, because they didn't attempt to make purchases with the stolen cards. They did, however, make veiled extortion threats by offering to furnish paid services that would "fix" the unpatched systems.

A Second Look

It's thus time for network professionals who gave up on the IDS a few years ago to go looking again. Indeed, market research numbers indicate that more and more of them plan to deploy IDSs in the coming years. Frost & Sullivan, for example, pre-dicts that the market for intrusion detection software will increase from $665.6 million in 2004 to $887.8 million in 2006 and $998.9 million in 2007. Another re-search house, IDC (*http://www.idc. com*), paints a slightly rosier picture, saying that the IDS market stands at $1 billion in 2005 and will grow to $5.6 billion by 2006.

Several developments have moved the IDS back into prominence. These in-clude IDSs' new ability to keep up with the high-speed transport technologies found in today's networks, the emergence of IDS "appliances," new intrusion-detection methods, better management tools, and a hybrid approach that com-bines the monitoring of the network- and host-based systems, the two basic types of IDSs, with a single console. The charge is led by many of the usual vendor sus-pects—Cisco Systems [5], Internet Security Systems (ISS), Intrusion.com, NFR Security, and Symantec—as well as numerous newcomers. The latter list includes CyberSafe, Entercept Security Technologies, and Enterasys Networks.

The market has also spawned a growing number of managed security services providers (MSSPs) with outsourced offerings that include intrusion detection ca-pabilities. In this area are Activis, Exodus Communications, OneSecure, NetSolve, RedSiren Technologies, Riptech, and Ubizen.

Moving to Anomoly Tracking

As noted, the developments driving the IDS marketplace are improving organiza-tions' ability to monitor and secure against unwanted attacks, whether intrusions or DoS/DDoS strikes. Arguably, the most critical is the growing use of anomaly-based intrusion detection by vendors of network-based IDSs.

The traditional network-based IDS discovers malicious traffic by detecting the presence of known patterns, a process usually called "signature matching." These systems work much like an anti-virus software package (detecting a known "bad" pattern generates an alarm) and effectively discover known patterns.

On the downside, signature-based network IDSs can suffer on two principal ac-counts. First, they can't see inside encrypted packets—the encryption essentially hides the packet's contents from the IDS, leaving it blind to assaults. Second, hack-ers often mutate the nature of their attacks, rendering pattern-matching useless. Just as an anti-virus package can't protect against a new virus until vendors patch their software, an IDS vendor must update its signature files—and it's not clear how many vendors have figured that out.

The anomaly-based network IDS uses packet sniffing to characterize and track network activities to differentiate between abnormal and normal network behavior.These devices analyze the data transfer among IP devices, permitting them to discern normal traffic from suspicious activity without pattern or signature matching.

These devices don't care about the content of data in a session (as with signature matching). They only care about how a session took place, where the connec-tion was made, at what time, and how rapidly (is a suspicious connection to one host followed by a suspicious connection to another host?).

With anomaly-based systems, it's important to get a baseline of what "normal" network traffic looks like. The chief difficulty of this approach is how to baseline— to know what's normal traffic as opposed to deviated. Signature-matching should be coupled with anomaly tracking. An anomaly can be compared against a signature, and if the anomaly doesn't show up on multiple probes, you ignore it. Cisco Systems, Enterasys Networks, Lancope, Intrusion.com, ISS, and Recourse Technologies are among the vendors that offer anomaly-based network IDS products.

Faster Systems

Most IDSs on the market now can keep up with a 400 Mbit/sec Ethernet. Beyond that, they begin to drop packets and become less efficient. When vendors push their IDS offerings beyond 400 Mbits/sec, they're only looking at a subset of packets. You can find products that will die in 400 Mbit/sec networks. Other players that boast IDSs capable of operating in 400 Mbit/sec network environments are Cisco and Enterasys.

Moving to Appliances

Another trend among IDS products is the network-based IDS appliance. Unlike first-generation IDS products, which required installing and configuring the ven-dor's intrusion-monitoring software on a PC, these appliances merge hardware and software into a preconfigured unit.

Cisco's Secure IDS, formerly known as the NetRanger, was among the first such appliances, and IDC believes this makes Cisco the current leader in this area. ISS (working with Nokia), Intrusion.com, and NFR Security (formerly Network Flight Recorder), are also moving their IDS products into the appliance category.

The appliance approach makes sense for several reasons. First, it eliminates many of the performance issues involved in installing IDS software on a general-purpose PC. The IDS software vendor can't optimize its product for every processor and revi-sion of operating system. Second, the appliance is a controlled environment, built to vendor specifications, so the IDS software can be configured specifically for the ap-plication. Appliance-based IDS boxes also eliminate operating system–related con-cerns, especially in all-Wintel or all-Unix organizations. Finally, appliance-based IDSs give plug-and-play capabilities to IT departments in multilocation companies and to service providers. These are especially valuable for deployment in remote offices, where novice end users can handle the physical connections while leaving setup and configuration to centralized IT staff.

IDS vendors have developed recent products that merge the capabilities of host- and network-based systems into a single management platform. In these environments, a management console works in conjunction with traffic- and log-analysis tools on the network and host IDS systems to provide a correlated view of network activity.

Correlating data from multiple network sources lowers the incidence of false pos-itives and enables network security personnel to view traffic from a higher level. For in-stance, a single scan of Port 80 on a Web server via a single router probably would not reveal the presence of an attack, but multiple scans across several routers would.

Outsourcing Intrusion Detection

Advances in IDS technology notwithstanding, organizations worried about unau-thorized intrusions and DoS attacks should also consider outsourcing their intru-sion detection needs. Outsourcing intrusion detection to an MSSP, which monitors customers' IDSs via the Internet, can make sense for several reasons. Not the least of these is cost. Companies with small, limited staff with limited experience in se-curity can benefit greatly from an MSSP. It would typically require five employees, working three eight-hour shifts (with extra staffing for vacations, sickness, and the like), to handle the 24-by-7 needs of an IDS-monitoring program. Forget about the $50,000 for the IDS—an employee costs at least $90,000 a year, and with five em-ployees, you could spend a fortune on training and maintaining security personnel.

Thus, it's important to sit down and perform a return on investment (ROI) study. During this process, IT organizations should ask themselves whether they have the expertise to operate critical systems that can cost a business revenue or customer confidence if they're compromised as the result of hacking or DoS or DDoS attacks.

The MSSPs tout the level of security expertise among their employees, claim-ing that this expertise enables them to better handle the task of deciphering often arcane IDS logs and alarms that befuddle typical IT employees. In addition, MSSPs have often deployed tools specifically designed to acquire and correlate information from a wide range of intrusion detection devices and systems. MSSPs Riptech and OneSecure, for example, both indicate that the technology they've developed in this area differentiates them from others in the market.

Riptech, for instance, spent two years developing proprietary data-mining and correlation software for its Caltarian security service. Caltarian's software permits the company to warn clients of attacks while they're under attack, with recom-mendations to protect their networks in real time.

So, perhaps one shouldn't pity the IDS after all. No longer an overly chatty box crying "wolf" too often, it now offers network managers an improved set of tools that can finally help them fend off unwanted attacks from insiders and outsiders alike.

Signs of Attempted and Successful Break-Ins

Hackers are succeeding more and more in gaining root-privilege control of gov-ernment computer systems containing sensitive information. Computers at many agencies are riddled with security weaknesses. When an attacker gets root privileges to a server, he or she essentially has the power to do anything that a systems ad-ministrator could do, from copying files to installing software or sniffer programs that can monitor the activities of end users.

The increase in the number of root compromises, DoS attacks, network recon-naissance activities, destructive viruses, and malicious code, coupled with the ad-vances in attack sophistication, pose a measurable threat to government systems.

In 2004, 599 systems at 43 federal agencies suffered root compromises in which intruders took full administrative control of the machines, according to the General Services Administration (GSA). That's up from 186 root compromises in 2002 and 332 in 2003. The government has only a vague idea of what kind of data may have fallen into the wrong hands.

For at least five of the root compromises, officials were able to verify that access had been obtained to sensitive information. For the remaining 594 incidents, com-promise of any or all information must be assumed. The compromised data in-volves scientific and environmental studies.

Meanwhile, the U.S. General Accounting Office (GAO), in a report recently re-leased, summarized security audits that have been completed at 35 federal agencies and indicated it had identified significant security weaknesses at each one. The shortcomings have placed an enormous amount of highly sensitive data at risk of inappropriate disclosure.

The government is going to find itself in "deep, deep trouble" if its IT security procedures aren't improved. If sensitive personal data about U.S. citizens is com-promised, Americans are going to wake up angrier then you can possibly imagine.

Many of the thousands of attempts to illegally access federal systems come from abroad. Also, many nations are developing information warfare capabilities as well as adapting cyber crime tools. Hackers exchange vulnerability information with one another. There is a whole new currency on the Internet that's called the back door. At-tackers trade information about back doors that provide access to different systems.

One step the government could take to increase the security of its systems is to focus more resources on improving education and training. Computer security ex-perts are scarce. They are in short supply, and they are expensive. The average salary is $120,000.

A 1998 directive by President Clinton, ordered all federal agencies to complete a virtual bulletproofing of their IT systems from attack by May 2005, but officials indicate that most agencies are behind in that work, and only a few are doing penetration testing.

Even more alarming, is that many attacks aren't detected. No one knows what was done, and no one has a way of knowing what was done.

Forensics

Threats to an enterprise's information infrastructure can come in a number of un-suspecting forms. Beyond fending off network intrusions and DoS attacks, compa-nies must stave off threats of industrial espionage.

Layoffs occur more frequently these days, and when the disgruntled, newly dis-enfranchised leave, today's technology makes it easy for them to sneak off with trade secrets, research materials, client lists, and proprietary software. Increasingly, cyberthieves are raiding corporate servers, electronically stealing intellectual prop-erty, and using email to harass fellow employees, putting companies at risk for lia-bility. The impact on the bottom line alone is cause for concern; the American Society of Industrial Security reports that theft of intellectual property in the United States costs businesses almost $6.9 billion annually.

Constant developments in information technology have posed challenges for those policing cyber crime. For many organizations, identifying, tracking, and prosecuting these threats has become a full-time job.

Specialists in computer forensics must use sophisticated software tools and spend enormous amounts of time to isolate anomalies and detect clues for evidence of a cyber crime or security breach. As previously explained, computer forensics is the equivalent of surveying a crime scene or performing an autopsy on a victim. Clues inadvertently left behind after a cyber crime can often be pieced back together to reveal details of wrongdoing and eventually pinpoint the perpetrator.

Although software tools can identify and document evidence, computer foren-sics is more than just technology and analysis. Safeguards and forensics method-ologies ensure that digital evidence is preserved to withstand judicial scrutiny and to support civil or criminal litigation should the matter be brought to trial.

Divining Good Forensics

Obtaining a good digital fingerprint of a perpetrator requires that steps be taken to preserve the electronic crime scene. The systematic search for evidence must adhere to basic guidelines to prevent the inadvertent corruption of original data during the course of investigation. Even booting up or shutting down a system runs the risk of losing or overwriting data in memory and temporary files.

The examination will usually begin with a look at the disk drive. Minimal han-dling preserves its integrity, so any disk investigation should begin by making a copy of the original, using the least intrusive manner available.

Today's forensic software tools can sniff out storage areas for data that may otherwise go unnoticed. Ambient system data, such as swap files and unallocated disk space, and file "slack" (data padded to the end of files), often hold interesting clues, including email histories, document fragments, Web browsing details, and computer usage time lines.

Be careful to document any inadvertent changes that may occur to the original drive data during data extraction. Complying with the rules of evidence preserva-tion and upholding the integrity of the process will help prevent any future chal-lenges of admissibility.

Although somewhat trickier than hard drive examination, data communi-cation analysis is another useful forensic tool. Data communication analysis typ-ically includes network intrusion detection, data preservation, and event reconstruction. Isolating suspicious network behavior also requires the use of specialized monitoring software. Doing so can reveal activities such as unautho-rized network access, malicious data-packet monitoring, and any remote system modifications.

Leave It to the Pros

Although today's sophisticated data-recovery tools have become fairly efficient, the process of recovery remains a tedious, labor-intensive task. And no matter how good the tools, the science of computer forensic discovery draws on multiple disciplines. Forensics demands a skill set often composed of software engineering and a solid familiarity with binary systems and memory usage, disk geometries, boot records, network systems, and data communications. Principals of cryptography are also important for identifying data encryption and password-protection schemes. Only experience can teach a forensic examiner how to avoid booby traps or an extortionist's logic bomb—items often left to wreak havoc along the path to discovery if not properly dismantled.

For these reasons, it's often wise to leave the process to the professionals. An expert in forensics will be able to quickly isolate the telltale signs of where to look for clues and will better understand data-discovery technologies as they apply to the legal process.

When selecting a forensic examiner, you should have several goals in mind: Your candidate should be familiar with the intricacies of your particular operating systems, know how to protect against data corruption and booby traps, and have a history of court appearances and controls established to deal with evidentiary pro-cedures, such as chain-of-custody.

If you're looking for more information on computer forensics or getting your staff trained on good procedure and practice, there are a number of good resources at your disposal. As storage capacities and network sizes continue to increase, so do the means by which cyberthieves can circumvent security as well as the effort re-quired to bring them to justice. So start training to detect the signs of suspicious ac-tivity today and learn how forensic computer investigation can protect your corporate assets in these dangerous times.

How a Hacker Works

Obviously, knowing how the hacker's mind works is only half of the battle. You must also know your network inside and out, identify its vulnerable points, and take the necessary steps to protect it. This section will look at some tips and tools administrators can use to prevent those vulnerabilities.

Diagram Your Network

You should begin by diagramming the topology of your network. You can do this with a sophisticated tool such as Visio, or you can use a less complex tool such as Word. Simpler yet, you can draw it by hand. Once you've diagrammed your net-work, identify all the machines that are connected to the Internet, including routers, switches, servers, and workstations. Then, evaluate the security precautions in place on those machines. You want to pay close attention to machines that have a public IP address on the Internet, because they're the ones that will be scanned by hackers.

Always-On Means Always-Vulnerable

Currently, the greatest security vulnerability is always-on Internet access using sta-tic IP addresses. With always-on access and a static IP, you are a like a big bull's-eye sitting on the Internet waiting to get hit. The question is, once hackers get in your network can they do any damage, or will they be frustrated and move on to the next target? If you have an always-on Internet connection, you should already have a basic security policy and firewall in place on your network. If you have a Web server, mail server, or other servers constantly connected to the Internet, your security responsibilities are even greater. Because the Internet is built upon the TCP/IP proto-col, many hacker attacks will seek to exploit the TCP ports of these servers with public IP addresses. A number of common ports are scanned and attacked:

Ways to Protect the Network

There are a number of ways to compensate for these vulnerabilities. First, you can implement firewall filtering. One of the best protections against port attacks is to implement a firewall with dynamic packet filtering, also called "stateful inspection firewalls." These firewalls open and close ports on an as-needed basis, rather than permanently leaving a port open where it can be identified by one of the hackers' port scans and then exploited. You can also analyze your system log files to track hacker activity. A third option is to install an intrusion-detection program that will do much of the log file examination for you.

Seeing What the Hacker Sees

In addition to protecting against the well-known vulnerabilities, you need to see what the hacker sees when he looks at your network. The best way to do this is to use nmap, a program that gives you a look at your network from a hacker-like perspective. A company called eEye has released a new version of this program for Windows NT (you can download it at *http://www.eeye.com/html/Research/Tools/nmapNT.html*). The company also offers an industrial-strength network security scanner called Retina, which helps discover and fix known and unknown vulnerabilities. This is an expensive, yet valuable, product.

Software Vulnerabilities

Hackers also often exploit software security problems. They take advantage of these behind-the-scenes parts of the software to gain access to your system. Thus, you should take stock of all the software running on your Internet-exposed systems. Go to the Web sites of the vendors that make each of the software packages and book-mark the page that has updates and patches for that software. You'll want to check these sites regularly and always keep your software up-to-date with the latest patches. Some companies even have services that will email you whenever there's a new update or patch.

Security Expert Web Sites

In addition to staying on top of your vendors' security updates and patches, you should also stay current on the security risks and problems that are identified by se-curity experts in the industry. Often, vulnerabilities may become known long before a vendor issues a patch. Therefore, your systems could be vulnerable for a period during which the hackers may know about it, but you don't. Two Web sites that will keep you informed are *http://www.atstake.com/* and *http://www.403-security.org*.

FINAL WORD: COMPUTER FORENSIC NEEDS AND CHALLENGES

Reporting of economic and cyber crime is problematic and grossly underestimated, as is apparent from the many risks associated with corporations' reporting or shar-ing fraud losses and activity. A uniform computer forensics crime reporting system should be developed that includes specific economic crimes.

The Fraud Identification Codes established by the National Fraud Center are a start. Until such a means of a computer forensics crime-reporting system is imple-mented and the stigma of fraud victimization is removed, this problem will not be solved. Uniform and thorough reporting is necessary in the war on economic and cyber crime; resources for computer forensics investigation and prosecution will naturally follow as the enormity of the problem unfolds.

The lack of agreed-on definitions regarding economic crime and computer crime has resulted in a paucity of data and information on the size and scope of the problem. Academics have not been able to agree on definitions and have, for the most part, continued to focus on white-collar crime.

Economic crime is defined as an illegal act (or a constantly evolving set of acts) generally committed by deception or misrepresentation (fraud) by someone (or a group) who has special professional or technical skills for the purposes of personal or organizational financial gain or to gain (or attempt to gain) an unfair advantage over another individual or entity. To this day, the true nature of the amount of eco-nomic crime is buried in the statistics of more conventional crimes. For example, credit-card fraud is typically classified as larceny instead of access-device fraud.

Preventing, detecting, investigating, and prosecuting economic crimes must become a priority in order to lessen their impact on the economy and the public's confidence. Law enforcement, as it stands now, is in danger of slipping further be-hind the highly sophisticated criminals. New resources, support for existing orga-nizations (the National Fraud Center, the National White Collar Crime Center, the IFC, and the Economic Crime Investigation Institute), and innovative computer forensics solutions are needed to control this growing problem in the United States and the world.

These computer forensics needs and challenges can be accomplished only with the cooperation of the private, public, and international sectors. All stakeholders must be more willing to exchange information on the effect economic and cyber crime has on them and the methods they are using to detect and prevent it.

No single sector holds all the computer forensics resources, tools, or solutions. In fact, industry has more resources than government, but it must be motivated and authorized to partner and communicate. All parties must be willing to work together to effect change in existing laws and regulations and to promulgate new initiatives. The victims need to follow the lead of the criminals and organize themselves, so that the organized bad guys are not operating in a lawless environment, where culpability is at a minimum.

The United States must take the lead. Current and future administrations must recognize the full impact of economic and cyber crime, both domestically and glob-ally, and make a concerted, strategic effort to combat it, for the benefit of all society.

Finally, let's move on to the real interactive part of this chapter: review questions and exercises. The answers and solutions by chapter can be found in Appendix E.

CASE STUDY 1: LOST FILES

A set of Word, Excel, and Project files that was created over 18 months relating to a project currently under construction has been maliciously deleted by a departing employee. The PC was not backed up. The action was discovered 3 days later and the IT group endeavored to locate and restore the files. They were unsuccessful. Management is assessing the options available. They are time consuming and ex-pensive. Some data cannot be rekeyed in because the source data is missing. The IT manager contacts a computer forensics firm. The firm finally restores the entire project directory within 4 days from first contact.

CASE STUDY 2: CORRUPTED FILES

Files relating to a multimillion tender on a sales and marketing PC have been found to be corrupted. The PC was not on the network and not backed up. The IT group advises that the data is gone forever. The tender closes at the end of the month, which is only 12 days away. Management is assessing the options available. The only option appears to be to withdraw from the tender process. Their hardware supplier recommends an inquiry to a computer forensics data-recovery firm. The firm re-ceives the hard disk at 4:00 P.M. on Friday and has a CD-ROM containing the draft tender response, worksheets, subcontractor quotations, graphics files, and periph-eral material on the client's premises by 11:00 A.M. on the following Monday.

CASE STUDY 3: DISAPPEARING FILES

The debtors module of an accounting package has somehow disappeared from the accounting PC. The software-support company is unable to locate the files, and the backup tapes do not restore correctly. The software-support company suggests that the data be rekeyed in—a massive task. Management is assessing their options. They are time-consuming and expensive. The distributor of the software recom-mends contact be made with a computer forensics firm. The firm finally restores the faulty data in time for the complete end-of-month statement run.

CASE STUDY 4: COMPUTER FORENSICS

The founder and majority shareholder of a consultancy business sold his interest to a multinational communications corporation. The contract of sale contained re-straint clauses, prohibitions on the removal of confidential information, and nonsolicitation of staff and client clauses. After about a year, the client—the multi-national—became suspicious that he was acting in breach of contract. A computer forensics firm was asked to investigate. At the outset, the firm suggested that the in-dividual's desktop and laptop computers be recovered to copy the hard disks and analyze their contents. Within an encrypted file on his desktop, the firm found a draft business plan for a new enterprise that would compete with his former busi-ness. On his laptop, in a deleted file that was restored, the firm recovered details of key clients and revenue streams. It was possible to demonstrate that information had been updated within these files after he had left the company, but before he had returned the computer. Taken together, the evidence was sufficient to initiate crim-inal proceedings.

CASE STUDY 5: FORENSIC ACCOUNTING

A multinational manufacturer reported significant losses in the company's distribu-tion division. It was not clear whether this was simply a result of an inequitable trans-fer pricing policy within the group or whether the company had been defrauded. Accountants from a computer forensics firm set out to investigate how the losses had been incurred, reconstructing incomplete records and unraveling a confusing series of transactions. They discovered that other companies within the group had trans-ferred products to the division at over market value to maintain their own profitabil-ity. More disturbingly, the division had sold much of its product at inexplicably low prices to a number of key customers. The business manager was dismissed after the computer forensics firm discovered that he had concealed ownership interests in some of these customers and evidence came to light indicating that he had accepted kickback payments. Poor and missing records prevented legal action from being commenced. In the following period, the division was on track to report profits fol-lowing tighter controls over transfer pricing and sales invoicing.

CASE STUDY 6: CORPORATE INVESTIGATION INTO PC PORNOGRAPHY

A computer forensics team was contracted to assist in an investigation for an orga-nization that suspected an employee of downloading and storing inappropriate ma-terial on a company PC. The team visited the site and, using correct forensic procedures, created an image of the hard drive of the suspect PC. The team was then able to recover a large amount of inappropriate material from the PC in a forensi-cally sound manner, including files that had been deleted, renamed, and hidden in an attempt to disguise their true nature. Using this evidence and the report the team produced, the client was able to take the appropriate action against the employee.

CASE STUDY 7: DATA RECOVERY

A computer forensics team was asked to assist an organization that had lost data as a result of a computer virus. The affected laptops were with field personnel and away from the central office when the virus was introduced. Consequently, the data collected over this period had not been backed up. The affected machines were brought to the team's secure laboratory, and, using forensic recovery techniques, they were able to image data from the affected machines, recover all of the data that had been stored since the machines had last been backed up, and eliminate the virus.

CASE STUDY 8: INDUSTRIAL ESPIONAGE

A computer forensics team was asked to assist in a case where it was suspected that industrial espionage had taken place through the computer system. It was sus-pected that a number of techniques had been used to plant spyware (remote con-trol and covert information-gathering programs) on a network. After carrying out a preliminary on-site analysis, the team removed a number of suspect machines to their secure laboratory for further analysis. A number of machines had been com-promised after employees had opened email attachments that contained trojan horse programs (programs that are disguised as common files but actually contain malicious code). Unfortunately, these had been missed by the organization's anti-virus measures. As an added service, the team's security engineers were able to offer advice and assistance in reconfiguring antivirus and firewall products to minimize the chance of a repeat occurrence.

CASE STUDY 9: FAMILY MEMBERS BOLT

Family members bolt, take the IT department and the product design, sabotage the originals, and go into competition. A family-owned product manufacturer and de-signer on the verge of being bought for many millions of dollars found most of its designs missing after the departure of key managers and designers. A program used for deep file destruction had been implemented to destroy both product designs and evidence of the procedure itself. An outside computer forensics consultant is brought in to recover designs and overwrites evidence instead. A computer foren-sics team is then brought in and discovers remnants of file destruction utility and data patterns consistent with sabotage by the same utility. The suspects finally ad-mitted to the use of the utility.

CASE STUDY 10: FORMER EMPLOYER

A former employer claims a competitor's new hire has stolen designs for manufac-ture. An individual working for a biomaterials firm gained employment with a competing firm. The individual had used several dozen diskettes for storage at the old firm and then used the same diskettes for new storage at the new firm. The pre-vious employer claimed that the individual took designs to the new employer on diskettes. A computer forensics team was engaged to demonstrate the employee's innocence. The original firm finally settled out of court.

CASE STUDY 11: GOODS LEFT TO ROT

Goods were left to rot while documents were allegedly backdated. A computer foren-sics team was hired to check the results of a police report that suggested the client's guilt. The client's attorney was advised as to the potential veracity of the claim. In-consistencies in the police report were discovered, and the sentence was mitigated.

CASE STUDY 12: MANAGERS START NEW COMPANY

Managers start a new company in the very offices of their employers; computers and backups disappear. A foreign branch of the entertainment arm of a multina-tional conglomerate suspects that key managers had been attempting to incorpo-rate company intellectual assets into a competing product line. Once the suspects believed they were under suspicion, the relevant office computers were reported as stolen. Data backups were reported as missing. Under pressure, the original com-puters were found and produced by a computer expert among the suspect group, but with large amounts of data missing. A computer forensics team was hired to in-vestigate. Unequivocal evidence of illegal activities was produced from the remains of files on the computers in question.

CASE STUDY 13: FAMILY MEMBER STEALS CLIENTS

A member of a family-run communications business left the company. While denying it, the individual started a business in direct competition with the family business. The individual's computer was identified as an asset of the original company. The individ-ual claimed that no company information was on the computer. A computer forensics team was hired to test the claim. Although the computer had been completely deleted, reformatted, and had entirely new operating systems and applications installed on it, the original database entries were, nonetheless, uncovered. The individual also claimed innocence up until the moment that the team experts were seen awaiting a call into the courtroom. The individual then admitted the wrongdoing and settled.

CASE STUDY 14: ERASED EMAIL

A private investigation firm was purchased, with a covenant by the previous own-ers not to compete. Within weeks, suspicion arose that the covenant was not being respected and that files and media that had been turned over had data removed. A computer forensics team was hired to look into the matter. Thousands of files were turned up by the investigation, showing a violation of the covenant.

CASE STUDY 15: BANK SUSPECTS

An employee of an FDIC-insured bank turned over a computer upon exiting from his employer. The managers suspected that this individual had revealed confiden-tial information regarding loan clients and credit information A computer foren-sics team was hired to inspect the email server records for deleted email files that might cast light on the individual's actions. In short order, the text of the suspect emails, which showed the former employee's culpability was revealed.

CASE STUDY 16: FORMER MANAGERS

Several managers left a software-design firm. Within a few weeks, they started up a new firm, producing similar products, in direct competition with the original firm. A computer forensics team was hired to inspect former managers' computers, which had been erased. Evidence that the business plan and designs for a new firm were taken directly from the original firm was uncovered. The new firm was enjoined by the court from offering their product until sufficient time had passed for them to have produced their own designs. The former managers were given a 9-month injunction.

CASE STUDY 17: FORMER CATALOG DESIGNERS

A company that had spent years producing a catalog and selling thousands of industry-specific parts found that a competing catalog with identical drawings and designs had been produced in a scant few weeks by a new competitor. A computer forensics team was hired to show that the designs of the new company were stolen from the original company. After the findings were presented to a court showing the original company's artwork and text being used in new catalogue, the new com-pany was enjoined from using designs for several months.

CASE STUDY 18: MODEL PURSUED

A wealthy suitor financed a young model. Once the model soured on the suitor, ev-idence was presented to show that the model had committed libel. Inspection of the

model's office computer was ordered. A computer forensics team was hired to at-tend the inspection only to find that the aforementioned suitor's computer forensic consultants did not understand how to access the data on the computer in question. With advice from the computer forensics team, an inspection was effected. With fur-ther advice from the team, it was shown that any suspect email was liable to have been forged by employees of the suitor. The suitor finally settled out of court.

CASE STUDY 19: ENCRYPTED MAIL

A former employee encrypted an email record, address book, and calendar to hide information from an employer. A computer forensics team was hired. The team then successfully cracked the file's encryption and revealed its contents.

CASE STUDY 20: TWO ATTORNEYS CAN'T SPEAK CIVILLY

Two attorneys couldn't speak civilly to each other. A computer forensics team was hired to act as a neutral expert when opposing attorneys did not trust results of each other's experts. The team brought the voice of reason to an acrimonious meeting between attorneys, and calm prevailed while the truth of the matter was revealed in the computer inspection.

CASE STUDY 21: BIG REAL ESTATE DEAL

The manager of a real estate fund was accused of increasing the value of shares in the fund by providing false information to potential investors when the value of the fund plunged. The manager was further accused of faking and falsely dating com-puter documents to support the claim of innocence. A computer forensics team was hired. Information that was recovered by the team mitigated and diminished claims against the client.

CASE STUDY 22: DOCTOR ACCUSED

A doctor was accused of withholding treatment based on the ethnicity of the ac-cuser. A computer forensics team was hired to inspect hospital records of treatment and meetings to support the medical provider's innocence. Deep inspection of all relevant computers and servers showed no evidence of wrongdoing.

CASE STUDY 23: FORMER EMPLOYEE CLAIMS

A former employee claims he never took any information with him when he left. The firm suspected the former employee of absconding with proprietary informa-tion. Under court order, the individual turned over a laptop computer, with no ob-vious data related to the case. A forensic inspection by a computer forensics team revealed enough relevant data to print two entire reams of documents. The suspect finally settled.

CASE STUDY 24: EX-PARTNER CLAIMS

A partner in an information technology firm left and went into his own business. The individual was accused of taking proprietary documents on his laptop. The individual produced the laptop, along with the claim that, although there were missing documents, none were relevant to the claims. Additionally, the individual claimed that a prolific virus had destroyed the documents. A computer forensics team was hired and was able to show fabrication of evidence, upon which the individual then admitted wrongdoing in a deposition. The individual was then sanctioned.

CASE STUDY 25: FORMER MANAGER

A manager of a Big 10 consulting firm went to work for a competitor. Under court order, the competitor provided a diskette that had gone with the individual to the new firm. A computer forensics team was hired to inspect said diskette. Although it was damaged, deleted, and overwritten, evidence of illegal customer lists and the lists themselves were discovered on the diskette.