PAUL MOBBS

A practical guide to sustainable IT

Unit 8

This unit is one of 12 sections to a "A practical guide to sustainable IT", a hands-on guide to working with everyday technology in an environmentally conscious way. The guide has been written by environmental activist and ICT expert Paul Mobbs, and was commissioned by the Association for Progressive Comunications (APC) with the support of the International Development Research Centre (IDRC). To download the full text of the guide, or any of the other units, please visit: greeningit.apc.org

A practical guide to sustainable IT

Author

Paul Mobbs

Copy-editing

Alan Finlay

Layout proofreading

Lori Nordstrom

Publication production

Karen Banks and Flavia Fascendini

Graphic design

Monocromo
info@monocromo.com.uy
Phone: +598 2 400 1685

Commissioned by the Association for Progressive Communications (APC).

APC

Conducted with support from the International Development Research Centre (IDRC).

IDRC | CRDI Canada

The views expressed in this publication are those of the individual authors and not necessarily those of APC or IDRC.

A practical guide to sustainable IT

Published by the Association for Progressive Communications (APC) with support from the International Development Research Centre (IDRC).

South Africa
2012

APC-201206-SU-R-EN-DIGITAL-162

ISBN: 978-92-95096-71-4

MAINTENANCE, SECURITY AND RESILIENCE

The term “sustainability” generally means the ability of a system to continue to function. In practice we must not only consider the ability to function, it's the ability to function even when external events are destabilising the system. Just like natural ecosystems, what a truly sustainable system must embody is not just the ability to work, but to continue to work after events have caused temporary interruptions or problems. This characteristic is called *resilience*,[1] and it's a feature that can be designed into the way we build and use computer technology.

This section looks at resilience of information systems in the broadest sense, considering how the equipment is used, how the hardware is maintained, and how the software is configured to maximise security. The objective of this process is to protect our information resources, and the machines that we use to access and manipulate them. That requires that we consider everything from how we define our needs and plan the system, the organisation of physical security, protection against malware and other system failures, right down to basic security measures such as passwords and access controls.

In many ways this section is all about combining all the actions covered in previous sections into a single, integrated method of managing our need for technology. Rather than separating different issues and looking at them in isolation we need to develop an integrated approach to system resilience and security, trying to minimise the ecological impacts of our demand for IT, and to refine the elements that make the system function to produce the optimal solution to our needs.

[1] Wikipedia, 'Resilience (ecology)'. en.wikipedia.org/wiki/Resilience_(ecology)

8.1. SUSTAINABLE DESIGN

There is no good or bad way to practice sustainable design[2] – it either works to serve your particular needs or it doesn't. In order to adapt to the diverse way in which we use ICTs it is often necessary to unpack and adapt certain ideas and strategies to fit your unique circumstances. The true test of any system is that when you depend on it to work at its best, it functions as designed to protect the integrity and reliability of the system and the tasks it is designed to perform – and does so while internalising ecological principles which minimise the impacts of the system upon the environment.

8.1.1. Examining options

There is no one method to achieve a good sustainable design. When we build in an ecological dimension to planning our IT needs it quickly becomes apparent that we have to change our working practices, not just the tools we use to carry out our work. For example, while using more efficient machines can lower energy use, changing the culture or expectation of the people using the equipment offers many more options to reduce impacts (for example, turning it off/using power-saving options when equipment is not in use).

The best way to approach sustainable design is to think more strategically about our present and future needs. Thinking over a longer period allows time for cultural changes/changes in working practices to evolve alongside technical change. In practical terms, what many organisations seek to achieve is a compromise between "deep green" sustainability, where ecological principles completely determine how we work, and more practical sustainability, where we try and implement what is possible today and set longer-term goals for improvement. How we do this is a matter of setting the context within which we ask questions or find solutions to problems:

- Take a long-term strategy which defines the need for certain equipment over a number of years – certainly longer than the lifetime of the equipment involved. By having an idea of what our needs might be in the future it is possible to invest in upgrades and replacements more strategically to reduce costs, and take advantage of new or more efficient systems.
- Always consider the effects of changing technical standards. Where certain technologies will obviously become obsolete we should plan for their replacement, although adopting new standards too early (for example, a new software package or operating system) might involve investment in equipment before its design has been perfected or its effectiveness proven.
- Technological standards can create inflexible restrictions – for example the use of proprietary designs which restrict compatibility with other similar systems (as a classic example, think of the many different types of USB or laptop power connectors in common use). Choosing hardware or software which uses generic standards avoids compatibility restrictions, and allows the switching of components/parts between different systems in the event of a fault in the equipment.
- It is important to question the cultural habits in the way people use ICT systems, just as we query the technical specifications of different technologies. Increasing efficiency and productivity often requires people to learn new skills, or adapt their methods of working to avoid habitual practices. Improving people's understanding and perceptions of the ecological impacts of their work, to be mindful of them, and to understand how they can work differently to address them, can be as cost-effective as finding technical solutions.
- Concentrate on what resources you already have. Optimising existing systems, using auditing or benchmarking tools to measure performance, can reduce the pre-existing ecological footprint – and might reduce the running costs or extend the working life of the equipment. More importantly, the process of optimising existing systems and working practices can highlight options or needs which

[2] Wikipedia, 'Sustainable design'. en.wikipedia.org/wiki/Sustainable_design

were not previously understood, and this information and experience can become essential when planning a longer-term strategy.

- Measurement is critical – “what is not measured is not managed”.[3] In order to compare the effect of two different ideas or options we have to be able to objectively measure their costs and impacts. Traditionally we use costs or prices to measure “the bottom line”. By building in values of both ecological and social impacts, in addition to traditional economic values, it enables decisions to be taken on a broader “triple bottom line”.[4]

Planning the transition towards a more sustainable way of working requires a balance between the ideal goals of sustainability, and the practical ability to find better methods of purchasing, operating and maintaining equipment. You must weigh up – as you perceive them – the risks, reliability and relative costs of different options to develop a sustainable solution to your IT needs, and then choose the set of options that can be reliably and easily maintained as part of your everyday routine.

8.1.2. Modular design

It can be difficult to consider the whole system design in terms of its sustainability. There are so many different aspects to sustainable IT that it is easy to become overwhelmed with the detail. To assist in the design process it helps to break down the system into its constituent parts, defining the larger system as a series of interconnected “modules”. For more complex systems, breaking down the problem into its constituent parts allows those with needs or expertise in different areas to evaluate the options that reflect their interest.

Creating a module design involves identifying the physical or logical boundaries which define each part of the whole system. The value of breaking down our needs or systems into smaller parts is that the problems related to each small part of the system are more easily quantified, and so can be more easily managed. Each part, independent of the rest, can be designed to function in the best way possible – and it is by systematically working through each module one-by-one that we progressively produce a more sustainable system.

For example, an ethernet network operates according to a technological standard. The machines which connect to the network use the same standard irrespective of which manufacturer's equipment is used. For that reason we can consider the network and the machines which connect to it as representing different modules, and each module can be evaluated as a single unit in order to select the optimum design. For network hardware, this division enables us to measure the different performance of a wired network versus wifi, or a single large network versus a series of smaller subnets, and then combine those options to produce a system with the least impact in terms of economic and ecological costs. The software that computers use to connect to a local server on the network also functions according to a technological standard, distinct from the network hardware. This allows us to consider the desktop machines people use as distinct from the servers, routers and other parts of the network – enabling measurements or research to be carried out to find the most effective/efficient solution for each of these parts.

By adopting a modular plan as part of a longer-term strategy for improvement it is possible to progressively upgrade different parts of the system in isolation, without necessitating the upgrading or modification of other parts. Organisational complexity also plays a role here. Breaking down the organisation's infrastructure into manageable sections which mirror the structure of the organisation allows roles to be assigned, projects developed and targets set to implement changes, and the monitoring of longer-term objectives to be reported.

For large IT systems, developing a modular plan can be an elaborate process involving consultation, planning and finding options which meet the needs of many different system users. For small or stand-alone IT systems it can be a more informal process because in practice there are fewer alternative options to consider. What's important in either case is that, for each part of the system, you have a clear view of the purpose and function of the module, the options available to provide it, and the costs and impacts of each option. The overall cost or per-

[3] Willcocks and Lester (1996). Beyond the IT Productivity Paradox, European Management Journal, vol.14 no.3 pp.279-290. dx.doi.org/10.1016/0263-2373(96)00007-2

[4] Wikipedia, 'Triple bottom line . en.wikipedia.org/wiki/Triple_bottom_line

formance of the system as a whole is then an aggregate of the results for all the modules – and providing that each module represents the optimum design, the whole system will then provide the best solution.

8.1.3. Building-in resilience

While modular designs allows more flexibility in planning or upgrading, making it easier to optimise performance, when we consider security and resilience we think of "layers". The concept of layering involves systematically building in security and reliability by spreading similar functions across different systems. This approach ensures that if one layer fails, then there are further layers to maintain system security and resilience. As with defining modules within a system, by breaking down the problem into individual layers, those responsible for that segment of the problem can take responsibility for it and work these tasks into their everyday routine – rather than having one person take on the whole task.

For example, the physical security in a building protects the computers inside the building; but if physical security fails and the computers are stolen then data back-up procedures should minimise the loss of data the machines contained. We can in turn break this example down into further layers to improve the system design. Building security can be broken down into separate zones, meaning that more important machines or information are protected by more levels of security. In a similar way, on a single computer non-sensitive information can be stored insecurely while secure data can be stored in encrypted formats.

Another common means of achieving better security and reliability is through redundancy.[5] In the event of one piece of equipment/a module in the system failing there is always a spare available to replace it. For certain applications where high reliability is essential (for example the safety systems of nuclear power stations) the same piece of equipment might be replicated two or three times – which of course has a much higher ecological impact. In high-availability computer systems multiple machines work in parallel, enabling the whole system to keep working if one or two machines fail, and even data storage might use an array of hard drives so that if one drive fails there is an exact copy of the data it contained on a second drive ready for use.

Unlike these highly specialised examples, improving the resilience of our everyday use of technology need not entail a large amount of expenditure and high ecological impacts. For example, on an average PC the only unique part of the machine is the data stored on the hard drive – the rest of the machine is expendable. By using a removable hard drive, taken from the machine at the end of each day and secured in a safe or locked cupboard, if the computer is stolen the data can be quickly reloaded from the old drive onto a new machine. Another option is to regularly back up the computer to an external hard drive so that, if the machine breaks down, the information on the external drive can be used on another machine.

Another way to develop resilience is to plan the purchase of equipment so that the likelihood of many failures occurring at the same time is minimised. All hardware has an expected working life – defined as a minimum by the guarantee period offered by the manufacturer. When a number of identical computers are purchased at the same time the chances are that they will begin to break down at roughly the same time too – especially if a certain batch of equipment contained a systemic flaw. Alternately, if you plan the procurement of equipment to a longer-term plan, then the purchasing or renewal of equipment can be staggered, and the chance that a large amount of equipment would break down over a short period is reduced.

At its simplest, creating resilience is all about having pre-planned alternatives to "business as usual". It begins with everyday procedures such as backing up data, so that if a computer fails you can still have access to the data it contained. In addition you might choose to have email accounts with more than one service provider, so that if one provider's system fails you can still send and receive messages. As part of the design process, when evaluating options for different elements within the IT system it is wise to provide alternatives to essential parts or services to maintain the system in the event of unexpected occurrences. If you have the in-house skills, failures can be quickly fixed or worked around by repairing or

[5] Wikipedia, 'Redundancy (engineering)

reconfiguring existing systems – perhaps utilising components from an expendable machine until longer-term repairs are made. For external services and support, it's important to have a regularly updated list of service providers or trades people so that in the event of a problem occurring with the regular provider you can quickly switch to another.

8.2. PHYSICAL SECURITY

The physical security of the room or building where equipment is kept might not appear to be a critical component in sustainable IT. However, IT systems and information appliances (digital cameras, music players, etc.) are expensive pieces of equipment; replacing them entails the expenditure of a large amount of energy and resources too. Ensuring that they are physically secure, both from theft and casual damage, is an essential part of making the equipment function for as long as possible, and keeping the information resources on the equipment secure. Improving the physical security around computer equipment must therefore be considered an essential part of how these systems are used.

8.2.1. The building

Addressing the security of a building is entirely dependent upon those who look after the building. If you control the building, that's simple – *it's your problem*. If the building is shared between many tenants that's a different problem. Likewise, if you share a single space, be that a formal office or a space in a community centre, you are reliant on others to do things for you to assist your security. In that case improving security isn't just your problem, it's about convincing others of the benefits they might get from the process.

The physical security of a building is reliant upon the individual elements from which it is constructed:

- *Doors and walls*. Doors[6] are a weak point because they are designed to open. A door can be strengthened by adding more hinges and locks, and reinforcing the door with extra wood or metal. If you use a deadlock, once locked the door can't be opened from the inside without a key – making it harder to remove objects if entry is gained by other means. However, there's no point strengthening a door to a point where it's much stronger than the walls around it (for example, a door set in a stud work and plasterboard wall) – otherwise it's just as easy to go in through the wall.
- *Windows*. Windows are a weak point, but often a last resort for illicit access due to the hazards of climbing over broken glass. Using key locks on window frames help because, once broken, the window frame can't be opened – which makes it much harder to climb through. The only effective way to secure a window is with internal bars or welded mesh securely fixed to the wall (external bars/fittings can be removed more easily).
- *Roof and floor spaces.* These are often overlooked. For adjoining buildings, if the roof space is shared then you'll need to secure any access point into the roof space from the inside. In offices which share a void above a suspended ceiling, you should also consider the likelihood of access from adjacent rooms. Likewise floor spaces can be vulnerable if their construction allows access from other rooms/spaces within the building.

While all the above are sensible measures, there is one very big note of caution. Yes, you can secure the building against various forms of external intrusion, either natural, accidental or deliberate. However, if the house/building is on fire and those inside can't get out, or the fire service can't get in, you've got a really big problem. The general point about security is that by working in layers you can spread the security risks through many different mechanisms rather than relying

[6] Wikipedia, 'Door security'. en.wikipedia.org/wiki/Door_security

on just one or two options. By spreading the security measures to secure the most important items, it can become unnecessary to put intrusive or expensive general security measures in place around the whole building.

Traditionally the way to secure a building is with a lock. Big heavy locks might give the appearance of security, but many locks have basic flaws which allow the perception of security they give to be quickly bypassed. Both the building, rooms, and cabinets/cupboards in the room can be locked, but don't rely upon mechanical locks to guard against access by skilled operatives. All key-based locks can be picked[7] if the person has the skills and the tools to do so. The standard front door lock, which clicks shut when pushed, is extremely insecure. By pushing a plastic card or a thin metal sheet (called a "shim") around the edge of the door it is possible to release the bolt and the door will open[8] – unless the mechanism has a dead lock.[9] Combination locks are more secure because they don't have a "key hole" that can be tampered with – but they still shouldn't be relied upon. Padlocks are also insecure, whether you use a key or combination version, because like front door locks they can be opened with a shim.[10]

8.2.2. The room

Working in one space makes it easier to secure the information and equipment that the space contains. If security measures can concentrate on that space, lesser physical security measures can be applied around the rest of the building. Work carried out in more than one location multiplies the security problem. General building security needs to be improved, and secure storage points – such as a filing cabinet or strong cupboards – need to be set up in each room/work location. You also have to give thought to how those measures are installed. For example, locked tamper-proof cupboards are not secure if they can be carried away – so if possible fix them to the wall or floor to prevent them being removed.

The greatest problem with securing the items in a room are those things which are physically difficult to secure – books, large desktop computers, DVD players, etc. Small high value items, such as digital cameras or external hard drives, can easily be locked in draws or cupboards. Laptops can also be locked away when not in use.

Which items we choose to secure in a room depends upon their value and/or whether they can easily be replaced. For those things which are replaceable – such as books, subscription publications and computer peripherals – the simplest option is to have insurance; if they're taken or destroyed you buy another with the insurance money (although some insurance policies might ask you to list all such items when obtaining the policy). If you want to keep other non-digital information safe – such as the original copies of important papers, certificates and other documents – then you'll need a fire-proof cupboard or safe to lock them inside when not in use. For irreplaceable items, such as film-based photographs or rare printed materials and books, the best option is to scan/digitise them, use the digital copies day-to-day and store the originals in a secure location away from the workplace.

The general problem is desktop computers, servers, laser printers and other expensive peripherals. If we look at where the "value" lies in computers, the hardware is expendable but the data that the computer contains is not – it's very valuable and often irreplaceable. For that reason it is easier to treat the hardware as expendable but organise the data it contains so it can be easily secured. On desktop computers use a removable drive caddy to hold the machine's hard drive so that, when not in use, the drive can be removed and securely locked away. Compared to the value of the information on the hard drive, a drive caddy is a minor expense, and can be easily installed in the machine. Laptops, mobile phones and other mobile devices should be backed up to some form of removable media – as outlined in unit 6.

8.2.3. Other building impacts

IT equipment uses electricity, but the environment within which these systems are used also consumes energy in the form of space heating, air conditioning, lighting, and fresh water for

[7] Wikipedia, 'Lock picking'. en.wikipedia.org/wiki/Lock_picking

[8] Wikihow, 'How to open a door with a credit card'. www.wikihow.com/Open-a-Door-with-a-Credit-Card

[9] Wikipedia, 'Dead bolt'. en.wikipedia.org/wiki/Dead_bolt

[10] Lock Pick Guide, 'Padlock shim'. www.lockpickguide.com/padlockshim.html

Figure 8.1.

Removable hard drive caddy

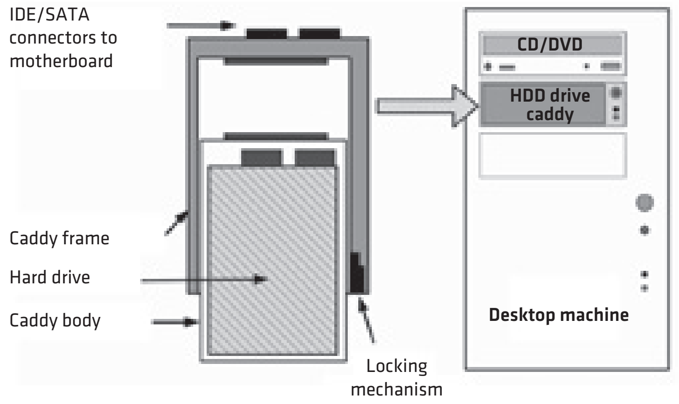

human consumption and flushing toilets. Creating a super-efficient computer system makes little sense if the environment within which it is used is not also optimised to minimise the impacts of the activities that the IT system supports. Again, this is the value of thinking in modules and layers – we can consider the needs of the whole system, and how best to address all the features involved.

When developing a sustainable IT system, it is possible to achieve similar, if not greater savings by attending to the design of the building and work areas.[11] This can be done through a formal building energy audit,[12] or less formal steps to tackle the main areas of energy and resource use:

- Space heating and/or cooling is the major consumer of energy in buildings,[13] often making up over half the energy budget of the building –

  - Space heating can be controlled through better insulation, but it is far more effective to reduce the operating temperature of the environment.

  - The need for cooling/air conditioning can be managed by reducing the heat load entering the workspace. CRT computer monitors use a lot of energy compared to flat screens, and switching to peripherals, printers and other devices which quickly switch to a low power standby mode when not in use will reduce the heat they produce.

  - A major source of heating is sunlight entering through windows – especially in more formal office environments which have large windows – simple light-coloured screens and curtains are the simplest way to reduce the solar gain from windows.

- Lighting is a significant factor in commercial/office environments. While compact fluorescent lighting reduces energy consumption significantly compared to incandescent bulbs, the latest tubular fluorescent luminaires and LED lighting modules[14] require even less power for the amount of light they can generate. Using timers on lighting systems can also ensure that lights are automatically switched off outside normal office hours.

- Water resources are an increasingly pressing ecological issue.[15] Public water supply requires one to two kilowatt-hours per person per day to treat and supply drinking water

[11] Whole Building Design Guide, 'Sustainable'. www.wbdg.org/design/sustainable.php

[12] Wikipedia, 'Energy audit'. en.wikipedia.org/wiki/Energy_audit

[13] Whole Building Design Guide, 'Optimize energy use'. www.wbdg.org/design/minimize_consumption.php

[14] Whole Building Design Guide, 'Energy efficient lighting'. www.wbdg.org/resources/efficientlighting.php

[15] Whole Building Design Guide, 'Protect and conserve water'. www.wbdg.org/design/conserve_water.php

– perhaps two or three times that figure in those areas which rely on desalinated water or water pumped for tens of miles from the nearest source. Minimising wastage and using low flush toilets are essential, and will become more so as the global climate warms.

- Indoor air quality is becoming a more significant health issue as research highlights the pollution created by modern construction materials and furnishings.[16] This can be exacerbated by sealed air-conditioned building designs which reduce ventilation. The solution is to use products which do not contain solvents, flame retardants, vinyl chloride and other volatile compounds which affect air quality.
- As well as using more sustainable IT systems, it's important to use products which have a lower environmental impact generally.[17] Using recycled materials helps, but as a general rule we should seek to use fewer goods overall, more reusable and repairable goods, and try to extend the working life of products for as long as possible.

One solution for electricity supply is to buy from a renewable energy supplier, or seek to develop a renewable power installation. This issue is considered in section 11.

8.3. THE MACHINE

Many aspects of the installation and configuration of the computer are dealt with in units 3, 4 and 5. All machines need care and maintenance – as outlined in Box 8.1 – and by regularly attending to the well-being of the computer system you can improve its performance and extend its operating life. What is equally important in the day-to-day use of these systems are the security-specific aspects of computer use.

8.3.1. Securing the machine

Computer security begins with the design of the "the box" – the case containing the components of the computer system. The most important component in your computer is not the CPU, or the memory, or even the motherboard – it's the hard drive(s). That's because the hard drive holds all the data on the system. Protecting the box is therefore primarily centred around protecting the hard drives and the data they contain. The simplest option is to install a hard drive caddy into the machine, as discussed earlier. The other option with desktop boxes is to get a cage or a clamp. These fit over the case and then bolt to a table-top or the floor. Commercial ones are quite expensive, but they're fairly easy to make by anyone with basic metalworking skills.

Laptops are a different problem. It is possible to take the hard drive (or, on new disc-less machines, the flash card) out of a laptop. Usually they slot into a compartment inside the body, accessed through a flap in the case. The problem is the connectors on the laptop's hard drive are not designed for regular removal; if you did this repeatedly they will become damaged. However, if you were in an exceptional situation where you felt it necessary to secure the data on the laptop, you could remove the hard drive from the laptop in order to store it in a more secure location.

For more effective laptop security the most common option is a cable lock. One end of the steel cable is bolted to a wall or worktop, and the other slots into the body of the laptop and is locked with a key, preventing it from being taken away. The problem is that the hard disk is easily removable – and if the sensitivity of the data on the machine is a concern you would need to use hard-disc encryption as well to prevent theft of/access to the data on the drive.

[16] Whole Building Design Guide, 'Enhance indoor environmental quality'. www.wbdg.org/design/ieq.php

[17] Whole Building Design Guide, 'Use environmentally preferable products'. www.wbdg.org/design/env_preferable_products.php

Box 8.1.

Care and maintenance of electronic equipment

The human world is not always kind to ICT equipment: Dust can clog cooling ducts and reduce heat dissipation; food and drink can be spilled on keyboards and other equipment; and for mobile devices, careless handling can slowly degrade the internal components until, perhaps after a very heavy knock, they fail.

All computers, peripherals and other consumer electronics hardware should have a section in their user manual on care of the device. This will give you the basics of how to care for your equipment, how to carry out any routine cleaning or servicing, and what methods not to use to clean the device. For example some video displays, especially touch screens, can be sensitive to the solvents and detergents used in many household cleaners. In general all electronic devices can be easily cleaned using an antistatic or microfibre cloth. Stains and splashes from food or drink are best removed by gentle scrubbing with a non-abrasive cloth using warm water with no added cleaning agents. If you spill a drink on a laptop it's usually easier to buy a replacement keyboard for the machine – and that's certainly cheaper than buying a whole new machine.

For mobile gadgets the greatest risk is dropping them on a hard surface. The simplest way to guard against this is to buy a purpose-built soft case to hold the device. In the event the gadget is dropped the case prevents a large mechanical shock being transmitted from the hard surface through the case. Many mobile devices, such as iPods or mobile phones, can be fitted with a lanyard so the device can be worn around the neck or wrist. Some even incorporate headphones so that you do not stress the headphone connector of mobile phones and MP3 players by tugging on the cord.

A major problem for digital electronics is static electricity. This can be generated by human-made fibres in clothes and carpets, and is made worse in a hot, dry environment such as an office. Static discharges can damage electronic components when the inside of the machine is exposed, or when you touch the metal contacts on cables or connectors. Another way of generating static electricity is from air moving across dry plastic surface. For this reason you should never use an ordinary household vacuum cleaner to remove dust from electrical equipment. Quite apart from the static risk, the high force of the vacuum can damage internal fans or rip the keys from keyboards.

If you wish to remove dust there are small low-suction vacuum cleaners available for use with electrical devices, which have conductive nozzles to prevent a build-up of static. You can also buy special anti-static brushes and cloths to remove dust without generating damaging static. Another option is to use small canisters of compressed air to blow the dust out of the case, fans and heatsinks. The cheapest option is to put your face close to the fan, keyboard or heatsink, close your eyes tightly to prevent damage from dust and grit, and then gently blow to remove the dust.

Dust is a particular problem with laptop computers. Laptops are often put down on top of soft furnishings and dusty table-tops – where the internal fan can suck up dirt and dust from beneath the machine and trap it inside against the heatsink. The internal heatsink of laptops use high density cooling fins with only a few millimetres gap between each one. Over a few months of inappropriate use these can easily clog with hair and dust, preventing the processor from being cooled adequately. You will notice this first when the area of the laptop where the processor is becomes unusually hot; then the processor will slow down during heavy operations because it can't lose heat at a sufficient rate. Finally the laptop may shutdown automatically to prevent damage to the system. Laptops usually have a panel in the base, sometimes secured by screws, which covers the fan and cooling fins – the user manual will usually show how to remove it. Then, using an anti-static brush or cloth, you can remove the wad of hair and dust which has built up between the fan and the heatsink.

Finally, as well as the physical hardware, the software systems of the device may need occasional maintenance. This is best carried out on a regular basis, such as regular calendar dates for following a routine back-up of information on the device. How this should be done is explained in the documentation written for the operating system, or on the producer's web site. An important part of system maintenance is installing software patches and upgrades. These not only fix system bugs, they also prevent malware infecting the machine from previously unknown flaws in the operating system. Some device manufacturers also offer firmware upgrades, to fix flaws within the programs which make the hardware function. Finally it is important to regularly “clean” the storage media inside the device of superfluous and unwanted files – for example, deleting the trash bin on the desktop. This is most appropriately done before performing a routine back-up.

8.3.2. Passwords and access controls

The purpose of access controls[18] is to prevent anyone who shouldn't have access to a machine, or a place, having access. We've had complex machine-based access controls in common use for a few hundred years – *keys*. Computer-based access controls are more complex, but that doesn't mean they're any less fallible than their mechanical counterparts. While there are recent biometric[19] or token access[20] systems available, the most common form of computer access control are passwords.[21]

Passwords are a means of access control, guarding against the unauthorised access to an information system. Some password systems are fairly weak, such as the PIN number[22] used with mobile phones and bank cards. Some passwords, such as the passphrases[23] used with data encryption, are stronger because they are more complex. The principle of password access is that, as there are so many possible alternatives, it's unlikely anyone could guess the password. How "strong" a password is relies on how many characters it contains, how many different symbols each character can represent, and whether it has a unique or predictable format.

To illustrate how strong these access controls are let's look at how they work. A PIN number usually has four digits, each with 10 possible alternatives (0 to 9). The number of potential PIN numbers is 10 × 10 × 10 × 10, or 10^4 (ten to the power four), or 10,000 possible alternatives. In contrast an eight-character password with at least 62 possible alternatives per character (A to Z, a to z and 0 to 9) has 62,^{8} or about 220-thousand-billion ($2.2×10^{14}$)possible passwords. Passphrases are usually much longer, and can contain punctuation characters and spaces. For an 18 character passphrase, with around 80 possible alternatives per character, that's 80,^{18} or eighteen million billion billion billion (or $1.8×10^{34}$)possible alternatives.

Manually entering a PIN number every 10 seconds, assuming an unlimited number of tries, it would take nearly 28 hours (10,000 × 10 seconds) to enter them all. This process is called a brute-force attack.[24] By using many processors in parallel, the specially designed "cracking" computers created for IT research and intelligence agencies can try *millions* of passwords *per second*. Against such professional cracking technology an eight-character password doesn't stand a chance and could be broken in a few weeks, but an 18-character passphrase might take months or years to crack.

In reality a brute-force attack doesn't start at "0" and step through each alternative. The process of password cracking[25] uses many steps to guess the most likely password/passphrase, based on the occurrence of certain words or letters in the language used. Another option is a dictionary attack[26] which tries words from a dictionary first, on the assumption that people use plain words for their password. By using people's habit of selecting whole words, names, birth dates, people's names and other such trends, the most common options can be eliminated first. However, an equally successful means of getting passwords is not the use of technology, but the use of human-to-human social engineering[27] techniques. Given the right scenarios, people will give away passwords, or can be tricked into doing so.[28]

There are two ways to create strong security using passwords:

Firstly, by using more complex and random passwords and passphrases. While there's all sorts of recommendations on the length and format of passwords – such as how many upper/lower-case letters and numbers you should use – research on the use of passwords suggest

[18] Wikipedia, 'Access control'. en.wikipedia.org/wiki/Access_control

[19] Wikipedia, 'Biometrics'. en.wikipedia.org/wiki/Biometrics

[20] Wikipedia, 'Security token'. en.wikipedia.org/wiki/Security_token

[21] Wikipedia, 'Password'. en.wikipedia.org/wiki/Password

[22] Wikipedia, 'Personal identification number'. en.wikipedia.org/wiki/Personal_identification_number

[23] Wikipedia, 'Passphrase'. en.wikipedia.org/wiki/Passphrase

[24] Wikipedia, 'Brute-force attack'. en.wikipedia.org/wiki/Brute-force_attack

[25] Wikipedia, 'Password cracking'. en.wikipedia.org/wiki/Password_cracking

[26] Wikipedia, 'Dictionary attach'. en.wikipedia.org/wiki/Dictionary_attack

[27] Wikipedia, 'Social engineering (security)'. en.wikipedia.org/wiki/Social_engineering_(security)

[28] Mitnick, Kevin (2003). The Art of Deception: Controlling the Human Element of Security. John Wiley & Sons, ISBN 9780-7645-4280-0 (paperback).

that this doesn't create much better security.[29] That's because people find it difficult to remember a truly random password, and so opt for easily guessable rules or formats which significantly reduce the variability and hence the security of their passwords. The solution is to use a mnemonic password,[30] created by a simple process which takes a more easily remembered phrase and reduces it to a string of letters and numbers (see Box 8.2).

Secondly, use multiple passwords for different functions on a machine. As more people use passwords on remote internet systems, the risks of using the same or similar passwords for all access controls is that if the passwords you use over a network are disclosed, it is possible to launch remote attacks on your computer system using similar style passwords to attempt a breach of security. Therefore when using multiple passwords, use different passwords, and a different style of password, on your local machine and on internet services. The difficulty is that remembering multiple passwords is difficult – but again, using mnemonic passwords based upon a more easily remembered phrase can solve this problem.

With most computer systems it is possible to set multiple levels of access control/passwords to secure access. First and foremost, it should be remembered that if the machine itself isn't secure – in particular the hard drives containing the data – then no amount of access controls will make the system secure. If an unencrypted hard drive is removed from one machine it can be plugged into another computer and read quite easily. Therefore, as outlined earlier, it is necessary to improve the security of all the layers in the system in order to make sure that security functions as an holistic process.

Assuming you've physically secured the system hardware, then there are various methods of using passwords to control system access:

- *BIOS password.* BIOS[31] is a firmware program held on the motherboard. If you set a user or system password, when the machine first starts up this has to be entered before the machine will boot an operating system. Unfortunately the BIOS password is easily circumvented by clearing the contents of the BIOS chip or removing the internal battery from the machine – although if that were the case you'd notice when you next used the computer because you wouldn't be asked for your usual password.

Figure 8.2.

Use of multiple passwords to secure desktop systems

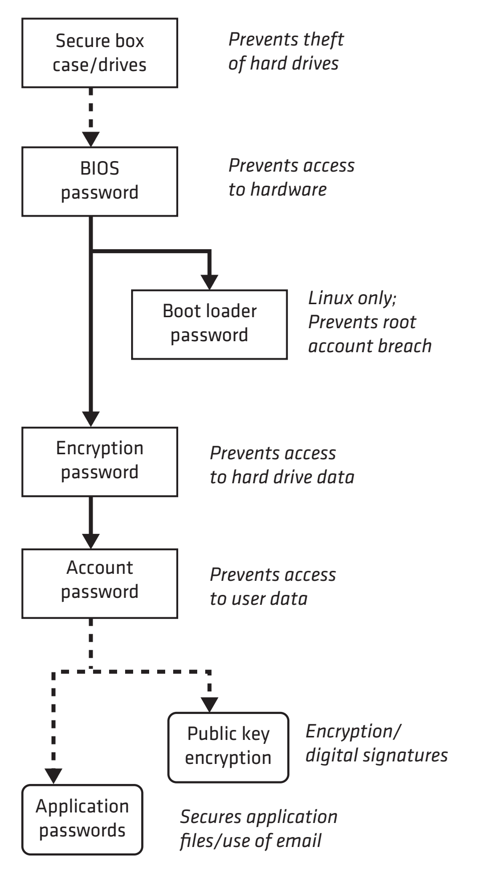
Secure box case/drives
Prevents theft of hard drives
BIOS password
Prevents access to hardware
Boot loader password
Linux only; Prevents root account breach
Encryption password
Prevents access to hard drive data
Account password
Prevents access to user data
Public key encryption
Encryption/ digital signatures
Application passwords
Secures application files/use of email

- *Boot loader password.* This is only of relevance on Linux-based operating systems. Ordinarily the boot loader program would load the operating system. When the machine boots it is possible to interrupt the boot loader to give it instructions – which can with a little knowledge of the Linux system be used to circum-

[29] Yan, Jianxin (2000). The memorability and security of passwords some empirical results. Computer Laboratory Technical Report 500, University of Cambridge. www.cl.cam.ac.uk/techreports/UCAM-CL-TR-500.pdf

[30] Computer Academic Underground (2007). Mnemonic Password Formulas: Remembering Secure Passwords. www.uninformed.org/?v=7&a=3&t=pdf

[31] Wikipedia, 'BIOS'. en.wikipedia.org/wiki/BIOS

Box 8.2.

Mnemonic passwords

Almost all user-based security features are based on the use of passwords, and understanding how passwords should be constructed and used is an important aspect of user security. For example, if the password “password” were used for many years, for every password on a computer system, that would be incredibly insecure; likewise if the same password was used with a sequential number added for different programs, that's insecure because the pattern can be easily guessed. In contrast, if each of the different passwords required to start and login to the machine were different, was only used for a few months or preferably less, and looked something like “cv6Td2Qb”, that presents a far greater security challenge.

When creating passwords people routinely substitute “1” for “i", “0” of “o”, “5” for “s”, or “3” for “e”, in order to add numbers to a dictionary word – creating something like “Pa5sw0rd” instead of “password”, or “acc3s5” instead of “access”. Such variations can easily by deduced and tested by password cracking programs, using a dictionary or word list to guess words. The greater problem is that because we need many different passwords, users might use just a few across the whole system, or they might add easily guessable changes – such as consecutive numbering of a root word (“password01”, “password02”, etc.).

To produce easily remembered and secure passwords you must find an approach that suits you. If you're the sort of person who can remember long strings of numbers and digits then you could use truly random passwords – using a random password generator to create them if necessary. If that's not the way you think then the strongest approach isn't to use a *mnemonic password* that reduces an easily remembered phrase into a string of characters. To make a mnemonic password begin with a favourite phrase, such as a line from a song or a poem. Then take the initial letter of each word in the phrase to make the password, while substituting easily remembered words or numbers to deal with repetition and significant meaning – such as substituting the word “space” with a space (' ') or underscore ('_') character, or the word “up” with the caret (“^”) character. For example:

- Take the phrase, “It is a far, far better thing that I do” – this reduces to “iaf,btId” (initial letters, ignoring the repetition of characters, and including the punctuation);
- The phrase, “may the force be with you” – this reduces to “mT4BwU” (alternates character case with each word, but uses the sound of “force” to indicate a number); and
- The word “encyclopaedia” – stringing out the word by its syllables this can reduce to “Ns1cl0PdA”.

By turning the characters of a password into “mnemonics” – small units that represent an idea – you can construct a seemingly random password in a way that's far easier to remember. We each have our own meanings and preferred ways of associating and breaking down words and phrases – which introduces some randomness in the setting of the mnemonic. As a result, even if the phrase is the same, two people may not reduce it to the same set of mnemonics vent the passwords of the system's "root" user account. Setting a password on the boot loader prevents this.

- *Hard-drive encryption password.* If the file system has been configured to encrypt the contents of one or more hard drives, you will be prompted for this password before the machine can finish loading the operating system. Hard-drive encryption is the best line of defence against the theft of the computer or its hard drives. The disadvantage is that it uses a lot of processor power to run the encryption process (a problem on older machines); and if you experience a data corruption problem you'll potentially lose access to all the data on that partition, or the entire hard drive.
- *User account passwords.* These are set when the computer is installed, and users can select their own passwords afterwards. The purpose of a user password is to protect data on a shared machine, or prevent opportunistic access to the user's files after they have logged in – for example, through the use of a password lock on a screen saver, and through the access permissions controlling access to files within the system. Although the system administrator can still change and override the user's account password, users can implement their own passwords outside the control of the system administrator using file locks and file encryption.

If you use a wireless or infra-red keyboard, that can be even more insecure because it's designed to transmit a weakly encrypted signal[32] which can be easily cracked by those with the required expertise and computing power.[33] In turn, any "strong" encryption codes entered and intercepted from a wireless keyboard can be disclosed by cracking the "weak" encryption its signals are encoded with. Basically, never use any wireless gadgets (including networks) for secure or sensitive work as they're a weak link in system security. Another potential problem is the local network. Any other computer connected to the network, due to deliberate use or infection with malware, can monitor data passing between local machines in order to detect passwords and other security information exchanged between machines.

8.3.3. File wrappers and locks

If you are the only user of a machine security is less of a problem; but if a machine is shared between a number of people that creates security issues that must be addressed. That not only includes desktop machines, but also access to a shared file or network server, and of course online services. As more data is lodged within "the cloud" using online services, how we secure the files we use from unauthorised access is becoming a more important issue.

There are two ways of securing the data inside files:

Firstly, using *file wrappers*. As the name implies, a "wrapper" is something the file is enclosed in in order to protect it. The two commonest forms of wrapper are file encryption and digital signatures. Wrappers can be used with files stored on the hard drive of the computer, files stored on a back-up storage media, and to protect files in transit on a network or files stored on remote systems. While the mathematical encryption of data is the basis of both file encryption and digital signatures, they serve very different purposes:

- *File encryption*[34] is intended to prevent access to the data in a file. Using mathematical functions the digital content of the file is scrambled to prevent access except by use of a digital key. In symmetric file encryption[35] both the person encoding the file and those decoding the file has to have a copy of the key – which presents a problem because sharing the key, by any means of communication, might lead to its disclosure. To avoid this, public key encryption[36] was developed. Anyone can encrypt data using the public half of the key and send it to the recipient. The recipient then decodes the data using their own private key – and only the holder of the private key can decrypt the data.

[32] Leyden, John (2007). Microsoft wireless keyboards crypto cracked. www.theregister.co.uk/2007/12/03/wireless_keyboard_crypto_cracked/

[33] Securiteam (2007). 27Mhz Wireless Keyboard Analysis Report. www.securiteam.com/securityreviews/6G0030KKKI.html

[34] Wikipedia, 'Encryption software'. en.wikipedia.org/wiki/Encryption_software

[35] Wikipedia, 'Symmetric-key algorithm'. en.wikipedia.org/wiki/Symmetric-key_algorithm

[36] Wikipedia, 'Public-key cryptography'. en.wikipedia.org/wiki/Public-key_cryptography

- *Digital signatures*[37] are not intended to prevent access to the content of the file, they are a means of protecting the contents of the file from being modified. As with encryption, a person uses a program to generate a digital signature, and this produces a block of data which can be attached to the file or held separately. Anyone can then check the authenticity of the data by checking the digital signature against the file using another program. With public key encryption this has become simpler as the public key can be used to verify the authenticity of any file that has been signed with the user's private key.

Many email programs now include features which allow the use of both digital signatures and encryption. This allows the easy encryption of emails or attachments to prevent disclosure; or by loading the public keys of email correspondents into your system, every time an email is received from them the program can verify that the email comes from that person (email is a very insecure system, and without digital signatures emails can easily be forged).

As shown in figure 8.2, each user can configure their own public- and symmetric-key encryption to protect their data on the machine, and use the passwords lock feature many applications use to secure files and the data they contain. These work separately from the security of the operating system – meaning that each user can implement their own data security irrespective of that applied at the system level. The principle security flaw in this process is the system administrator. They have the ability to log the user's activity using key-logging software which monitors the keystrokes entered from the keyboard. Therefore, if you don't know or trust the system administrator of the machine you have to work on, it's best to assume that the system is not secure – even if you have the ability to configure high-grade encryption of the data on the system.

The second type of password protection is the built-in passwords used for many popular office-based applications. This allows the use of passwords to control access to the content of files – without the password the file cannot be opened and read. Some applications, such as the programs used to create PDF files, allow you to set conditions for how the document can be accessed – preventing those accessing the file from printing, using cut-and-paste or modifying the contents unless the file is unlocked with the correct password.

The general problem with the file locks used by application programs is that they're very weak. Most applications – such as PDF files, word processor files and spreadsheets, or ZIP archive files – have downloadable "cracking" programs available to find the password and unlock the file. These are brute-force programs, and with a reasonable amount of parallel computing power it is possible to crack the passwords on applications within a matter of hours or days. In contrast, if you wish to protect the integrity of the document's contents rather than just restrict access to it, cryptographic digital signatures are far more complex and can't be easily circumvented.

[37] .Wikipedia, 'Digital signature'. en.wikipedia.org/wiki/Digital_signature

8.4. MALWARE AND PHISHING

Malware[38] – or to use the more popular label, “viruses” – have been a problem with computers for at least the past 20 years. Originally they were carried in files and on floppy disks swapped between machines. Today, with the advent of high-speed broadband, the malware problem has become much larger and more complex, and can be transferred from the active code used in websites, email, as well as software programs. The motivation for producing malware has also changed; 20 years ago writing malware was a past-time for computer geeks and pranksters, whereas today the production of malware is increasingly related to organised crime and the work of intelligence agencies.

Most of the viruses, trojans and related malware in circulation are designed specifically for use with the Windows operating system. That is partly a reflection of the widespread use of Windows – if you want to write malware for nefarious purposes then it makes sense to target the most popular operating system. Both Linux[39] and Mac OS have far fewer problems with malware. That's partly because they are used less than Windows, but it is also related to the design of the operating system itself. Unix-like operating systems, such as Linux and Mac OS, are designed to be more secure and to enforce strict control over how users can use the operating system, unlike the Windows system which has traditionally been designed for convenience and ease of use. There *are* examples of Linux and Mac viruses. In additional there are also cross-platform viruses that can use Java, or the scripting language used in office-based programs such as Microsoft Office and OpenOffice, which can affect all major operating systems. Even the Android mobile phone system now has malware developed for it.

Malware isn't just a security issue. Machines with malware running in the background use more processor power, and if part of botnets[40] they will use the broadband connection to shift large amounts of data as part of spam or denial of service attacks. This results in higher energy consumption, and for those whose broadband connection is metered it can also result in very large bills for the data sent over the network. However, it is also true that running anti-virus software, which routinely scans the operating system and incoming/outgoing data for malware, will add to the load on the processor and use more energy, and on older machines it might reduce system performance slightly.

For Windows machines anti-virus software is an essential part of running an internet-connected system. There are some free anti-virus packages available for Windows – most notably Microsoft's own *Windows Security Essentials.*[41] Programs which are paid by subscription usually provide a higher level of protection, providing not only regularly updated anti-virus functions but also the filtering of web traffic and email. A particular problem with Windows are trojan programs which infect the machine and monitor its operation, potentially sending secure personal data to online fraudsters who can use it to commit crime.

For Mac and Linux machines anti-virus software is still an optional extra as it's not essential to everyday operation – although it might be if these systems grow in popularity. Where some form of anti-virus software is required is on Linux servers; not to protect the server, but to prevent malware being transmitted by the server which might affect Windows users. While Linux users often talk of the system's immunity to malware, it is just as susceptible to unknown security flaws as other systems – although one feature of open source software is that programmers are free to study the code in order to find such flaws and fix them. Even if someone were to knowingly download a program or script containing

[38] Wikipedia, 'Malware'. en.wikipedia.org/wiki/Malware

[39] Wikipedia, 'Linux malware'. en.wikipedia.org/wiki/Linux_malware

[40] Wikipedia, 'Botnet'. en.wikipedia.org/wiki/Botnet

[41] Microsoft (accessed June 2012). Microsoft Windows Security Essentials. windows.microsoft.com/en-GB/windows/products/security-essentials

malware and execute it, depending upon the permissions used to configure the system, it is likely to affect only the user account it is run from. The use of the suite of security hardening tools which are now available for use with Linux, such as *Security-Enhanced Linux*,[42] reduce the likelihood that malware would cause significant damage to the system.

Lastly, an increasing problem on the internet is *phishing.*[43] Phishing is the use of the internet as a means to make people disclose sensitive financial and personal information. That information can then in turn be used to perform various forms of fraud and identify theft.[44] Often such frauds ask people to pay small sums of money as a "handling charge" in order to arrange the transfer, or to receive the goods promised, which ultimately will never turn up.

Email users will regularly receive emails telling them that they have won money, or someone wants to give them money, in return for their bank details. The best response to emails such as this is to hit the delete key. A significant problem is the use of HTML emails as this allows the display of web addresses to indicate one web location, such as a bank, when the underlying link takes the user to a web site where they will be persuaded to release sensitive information, or load software which might install malware on their system. Before clicking on any link in an HTML email it is always essential to look at the web address. Often this will give an indication of fraudulent use as the address does not match the site/organisation it claims to represent. While a more technical issue for novice users, if you inspect the header data contained in the email it is possible to look at which address the email originated, and then check if that address matches the real address of the organisation's claimed identity.

Another recent phishing phenomena are hoax phone calls where an operator tells you that they have "detected a problem with your Windows computer"[45] (at present this is a problem specific to Windows – Mac/Linux users usually have the call terminated when they say that they do not use Windows). The caller asks

[42] Wikipedia, 'Security-Enhanced Linux'. en.wikipedia.org/wiki/Security-Enhanced_Linux

[43] Wikipedia, 'Phishing'. en.wikipedia.org/wiki/Phishing

[44] Wikipedia, 'Identity theft'. en.wikipedia.org/wiki/Identity_theft

[45] Microsoft (accessed June 2012). Avoid tech support phone scams.
www.microsoft.com/en-gb/security/online-privacy/avoid-phone-scams.aspx

you to perform some commands on your computer, and then pay to download some software to solve the problem. In fact, apart from the scam of making you pay for software that is not required, downloading that software is likely to introduce real malware onto the system.

As a general anti-phishing/online fraud avoidance rule, any organisation wishing to offer money, products or IT support, either on the phone or online, should not object to giving you a company name, land-line telephone number and postal address at which you can contact them. In order to ensure that you have redress for any consumer fraud, it's important that the organisation is based within your resident legal jurisdiction. Any organisation which is not willing to give a telephone number or postal address to verify their identity, or which insists on conducting business only through websites or email, should not be trusted.

A practical guide to sustainable IT

This practical guide to sustainable IT offers a detailed, hands-on introduction to thinking about sustainable computing holistically; starting with the choices you make when buying technology, the software and peripherals you use, through to how you store and work with information, manage your security, save power, and maintain and dispose of your old hardware. Suggestions and advice for policy makers are also included, along with some practical tips for internet service providers.

Written by IT expert and environmentalist Paul Mobbs, the purpose of the guide is to encourage ICT-for-development (ICTD) practitioners to begin using technology in an environmentally sound way. But its usefulness extends beyond this to everyday consumers of technology, whether in the home or office environment. We can all play our part, and the practice of sustainable computing will go a long way in helping to tackle the environmental crisis facing our planet.

This is also more than just a "how to" guide. Mobbs brings his specific perspective to the topic of sustainable IT, and the practical lessons learned here suggest a bigger picture of how we, as humans, need to live and interact in order to secure our future.

The guide is divided into 12 sections (or "units"), with each unit building thematically on the ones that have come before. They can be read consecutively, or separately. The "unit" approach allows the sections to be updated over time, extracted for use as resource guides in workshops, or shared easily with colleagues and friends.

The guide has been developed on behalf of the Association for Progressive Communications (APC), with funding support from the International Development Research Centre (www.idrc.ca). It is part of a APC's GreeningIT initiative, which looks to promote an environmental consciousness amongst civil society groups using ICTs, and amongst the public generally. Other publications and research reports completed as part of the GreeningIT initiative can be downloaded at: greeningit.apc.org

APC

IDRC | CRDI