Hardening Your Computing Assets †

Carlo Kopp ††

Carlo.Kopp@aus.net

ABSTRACT

Computing equipment designed to commercial standards is susceptible to a wide range of electromagnetic attack techniques. This introductory paper reviews the susceptibilities of such equipment and proposes some measures for hardening at a site level, and hardening equipment by design.

Computing equipment designed to commercial standards is susceptible to a wide range of electromagnetic attack techniques. This introductory paper reviews the susceptibilities of such equipment and proposes some measures for hardening at a site level, and hardening equipment by design.

1. Defending Against IW

The first question which we must ask is whether we are at risk or not, and what are the likely losses to be incurred should we be successfully attacked. This provides baseline for the budget to be allocated to the task of protecting our site. The second question is then that of what is the most likely mode of attack.

Cyberwar attack or hacking will in many instances be the preferred mode of attack, but in some instances electromagnetic attack intended to cause denial of service for short or long periods of time may be a possibility. In the short term, electromagnetic attack is not particularly likely, although some reports from Europe and the US suggest that it is beginning to occur. Once police forces worldwide start deploying HERF guns for traffic control purposes (see an early November issue of New Scientist for more detail here), the technology will however become more available, thus better understood in the wider community, and the frequency of incidents will inevitably increase. The law enforcement community should give some careful thought to the fact that in promoting the proliferation of the HERF gun to solve one law enforcement problem, they may have inadvertently opened a Pandora's box of other law enforcement problems, potentially far more expensive to the general public.

Having determined that we are at risk from electromagnetic attack, we must then determine what the likely style of attack will be. The threat can be divided into high power and low power styles of attack. High power attack, by flux generator bomb or microwave bomb, is less likely but considerably more damaging. It is less likely because the technology is difficult to produce without the resources of a government research establishment, and the equipment to perform this kind of attack requires often difficult to source materials, such as high grade plastic explosives, high performance detonation systems similar to those used in nuclear weapons, and finally a non-trivial amount of expertise is required to use these weapons properly. Delivery may also prove to be an issue, as a high power flux generator requires a packaging volume similar to that of a sizeable car bomb. High power attack is therefore only likely in the instance of war, or a terrorist attack sponsored by a hostile government prepared to provide the logistical support for the weapons. It is worth noting that any government with the ability to build an implosion type nuclear bomb will have the

---

† First published in the February, 1997, issue of Asia/Pacific Open Systems Review, Computer Magazine Group, NSW, Australia under the title of "Information Warfare - Part 2".

†† Carlo Kopp, MSc(Comp.Sci), BE(hons), is a former computer design engineer, embedded programmer, Unix systems programmer and a practicising Unix systems consultant, with over 15 years of industry experience. At this time he is working on his PhD in Computer Science at Monash University. He may be reached via Carlo.Kopp@aus.net or http://www.cs.monash.edu.au/~carlo .

required hydro-dynamics expertise to eventually design themselves a flux generator or microwave bomb.

Low power attack is more likely simply because the expertise required to build a HERF gun is much lower, the components are quite readily available, and finally with the expected wide proliferation of police HERF guns these will be relatively easy to acquire. Since they are not as yet covered by legislation in most countries, it is entirely feasible that such "Luddite specials" will be smuggled across national boundaries.

Other modes of low power attack, such as disruption of mains power supplies and Tazer attack on LANs, are also more probable because the technology to do these is readily available, particularly in the US.

There are a wide number of measures which can be applied to hardening a site against such attacks, and we will now take a closer look at some of the alternatives.

2. Site Hardening

Site hardening is based upon the model of electromagnetically "soft" computer equipment being protected from exposure to damaging voltages and electromagnetic fields. In this fashion, the user has the choice of arbitrary computer equipment, which is survivable because it is never exposed. The weakness of a site hardening approach is that an attacker who can penetrate the site's security perimeter, the hard protective shell of the installation, may have the opportunity to do much damage.

The starting point for site hardening must be networking, as networking cables are exceptionally good at propagating damaging voltages which may be coupled into them, and because networking interfaces are often not designed to handle anything beyond trivial levels of power.

Fortuitously, the network is probably the easiest part of a site to harden, because optical fibre variants of most networking interfaces are readily available. Indeed, the 10-Base-T and 100-Base-T standards support optical fibre versions, which are commercially available off the shelf. Whilst they may be slightly more expensive than their copper equivalents, they are wholly immune to any kind of electromagnetic attack, as well as being inherently immune to problems with building earths, and lightning strikes. Because a short run of a 62.5 graded index optical fibre cable also has a bandwidth running into Gigabits/sec, it is also a worthwhile investment in the long term, unlike 100-Base-T and ATM twisted pair cables which will cease to be useful in a decade or so, when the industry moves up to Gigabit/sec network speeds. Needless to say, while these are all excellent reasons for installing optical fibre LANs, rather than copper LANs, the short term cost overheads, and a lack of familiarity in the marketplace mean that to date, only a small proportion of the existing LAN base has adopted this superb technology.

The next item on our list is the protection of the mains power supply. Just as LANs represent a potential single point of failure for many machines on a site, so does the mains power distribution in a site. High voltage spikes or high RF voltages injected into the mains power will find their way to almost every machine hanging off the mains. Because most machines employ low cost volume production switchmode supplies, particularly PCs, it is possible that damage may be caused directly to the power supply. Should the power supply cope, then it is entirely possible that damaging voltages can couple through the supply into the equipment to damage or disrupt logic devices. Surge suppressors and uninterruptible power supplies (UPS) may be quite ineffective, as most of these are designed to cope with much less destructive circumstances. Of major concern is that many UPS types are of the cheaper standby/bypass variety, which route mains power directly through and couple the standby battery power in only if the mains goes down. As a result, a large spike or RF voltage injected into the mains will pass through them quite unhindered. In any event, a low cost fully isolated UPS may not survive an RF or spike hit coming in from outside, leaving the site with X minutes of uptime before the batteries are exhausted. If the UPS is bypassed to get the site up and running, the next hit will get at the site's computers directly.

As with networking, a simple and relatively cheap measure may be used to solve this problem entirely. And it is not by any means a new idea. The solution is the use of a motor-generator power isolator. Such devices are simply built by coupling a robust single or triple phase electrical motor to a single or triple phase generator (alternator). Mains voltage powers the generator, which produces clean mains power for internal distribution within the site. It is worth noting that motor-generator power converters were commonly used in the early days of mainframes, as the then dominant linear power supplies did not cope well with poor quality mains power and thus motor-generators were a must to achieve good uptime and avoid

mainframe crashes.

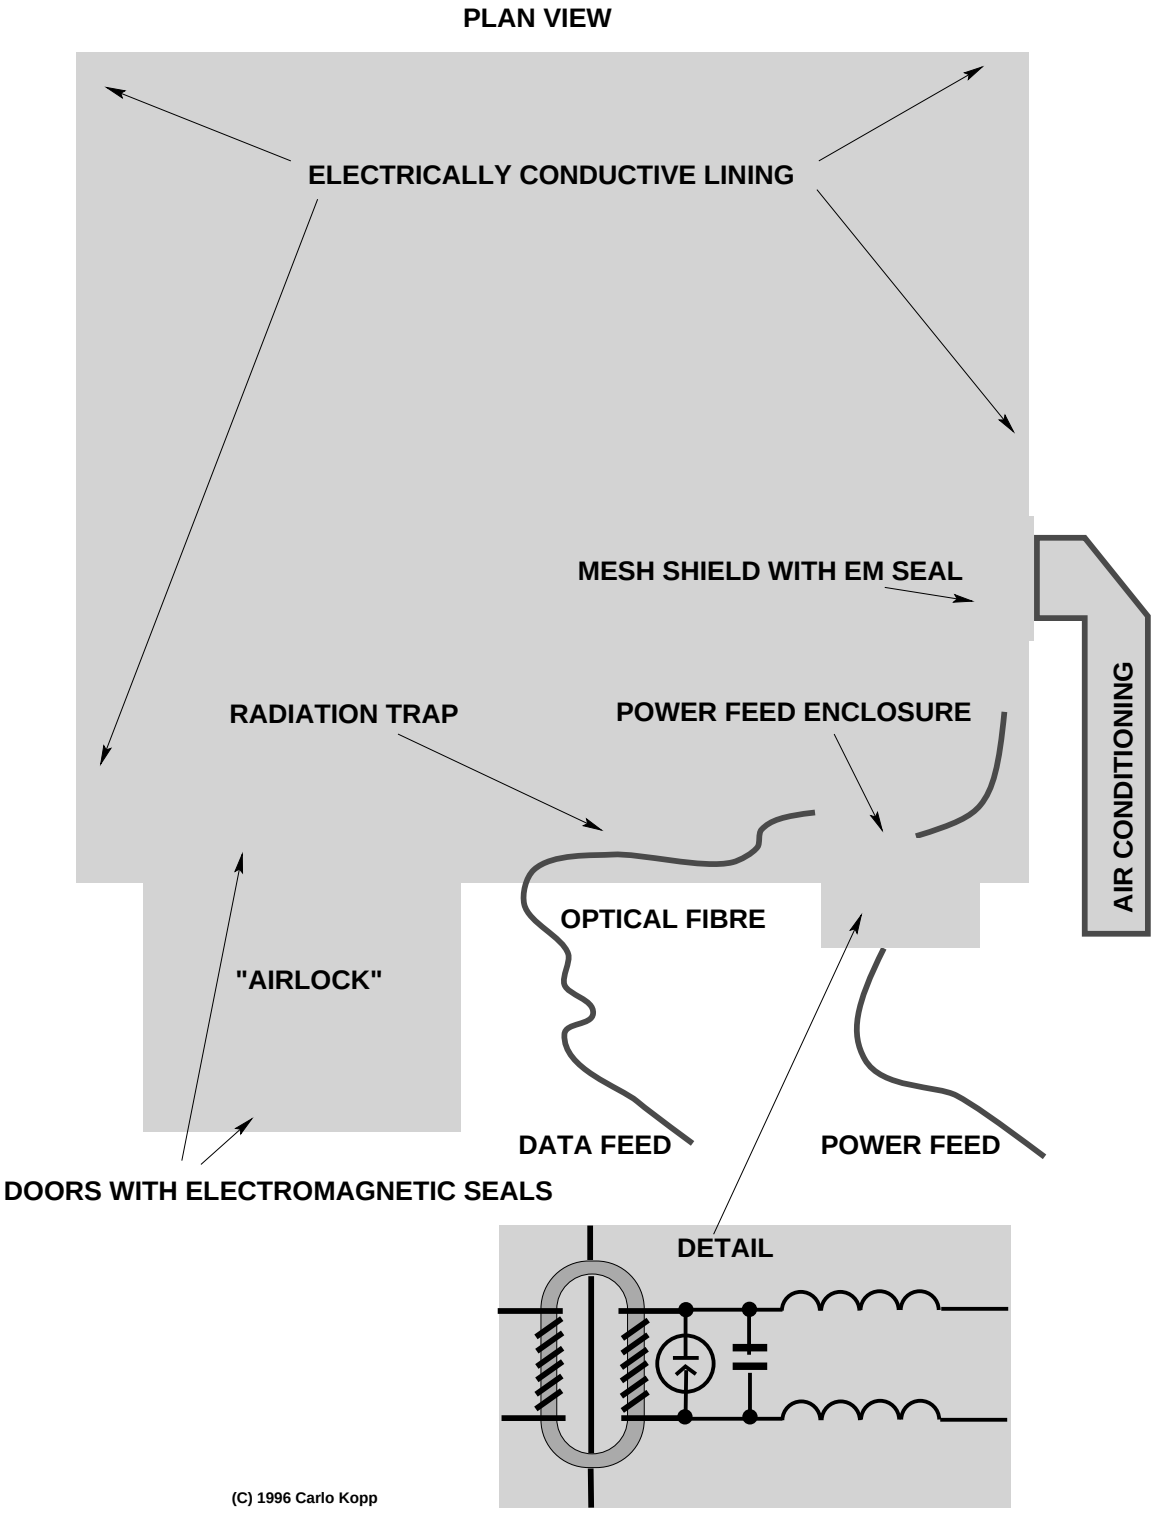

So, if you have an old motor-generator mains converter in the company warehouse or basement, now may be the time to dust it off and give it a new lease of life. If not, you may have to expend a few thousand

dollars and get one built or ordered. Needless to say, the motor-generator should be installed upstream of the UPS or any other lesser protection devices within the site.

An important note here is that a motor-generator will protect the site from disruption injected into the mains from outside the building. A HERF gun aimed into the building may couple directly into the mains wiring <I>downstream</I> of the protection device.

If we have installed optical networking and protected the building or computer room power supply, we have made life very difficult for an attacker. He will have no choice than to specifically target a given room or area in the building with his HERF gun to couple into local mains wiring, keyboard/mouse cables or SCSI cabling, or get through the shielding of the machine in question. However, a number of straightforward solutions exist to this mode of attack as well.

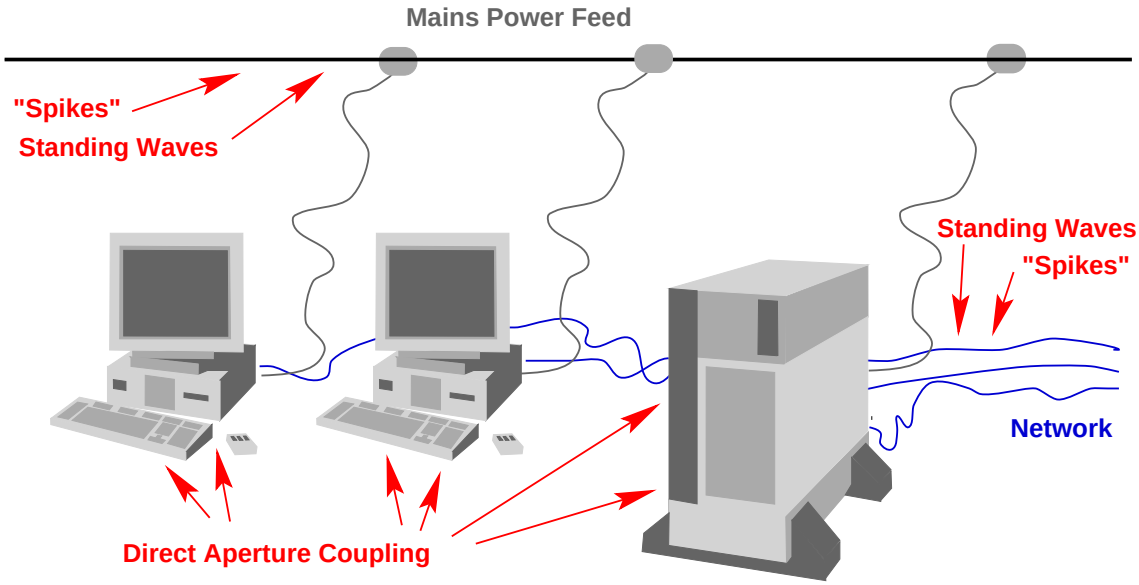

Fig.1 System Level Susceptibility

To solve this problem we can use Faraday cage or electrostatic shielding to simply exclude RF signals from the environment occupied by the equipment. A good measure of protection may be applied to mains wiring by running it through metal enclosed cable trays, and threading it through flexible metal shielding armour from the tray to the wall socket where it is to be used. A wide range of off-the-shelf commercial products may be used to this end.

A computer room, office or even equipment cupboard may also be built or refitted as a Faraday cage, by covering the walls, floors and ceiling, windows and doors with conductive copper mesh. Even a very basic shielding arrangement will exclude much of an impinging RF signal. For the very security conscious, this will also disable bugs and mobile telephones.

If the site is critical, then comprehensive Faraday cage shielding may be warranted. Too do it thoroughly, the shielding must not only be comprehensive but even small gaps and apertures will need to be thoroughly sealed. This means that incoming and outgoing cables will need to be routed through RF traps or Ferrite grommets, doors, windows and air conditioning vents will need proper flexible seals (similar to that in many microwave ovens), and a airlock arrangement may be needed for the door. Phone lines in and out will also need to be coupled through optical fibres. Needless to say, this can get to be quite expensive, particularly if the intent is to shield a large area or whole building. However, all of the required technology has been available off the shelf for many years, due the US military EMP hardening effort during the Cold War. The author has a 2 inch thick stack of component, seal and material catalogues, most acquired locally.

As is quite clear, the means of defending against most types of electromagnetic attack have been around for decades, and provide if applied properly, an excellent measure of security against nearly all threats lesser than a flux generator bomb on your site doorstep. Players in the latter league will be more

inclined to park a van loaded with ammonium nitrate in your site basement.

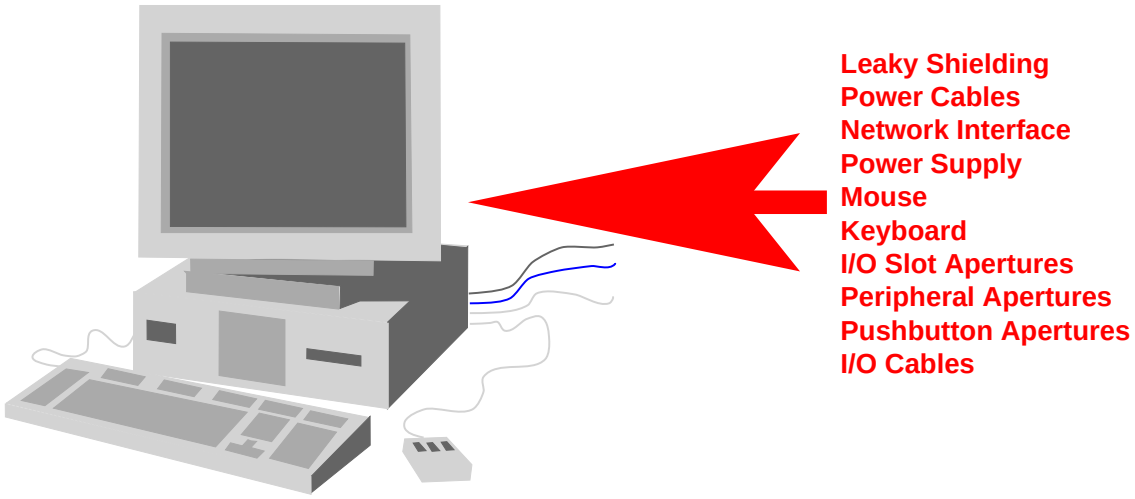

Fig.2 Host Level Susceptibility

3. Hardening Computer Equipment

Hardening of computer equipment by design would obviate many of the cost overheads which may be incurred through site hardening. Indeed, if a piece of computer equipment is sufficiently robust, then it may not be necessary to apply any hardening measures to the site as a whole. However, hardening by design requires that the vendors of the equipment cooperate and since they are largely oriented toward to minimising production costs of equipment to maximise margins, only intense pressure from the customer base will produce the desired effect. Arguably, should a sufficiently large number of corporate and government customers demand robust equipment, then we may see some results.

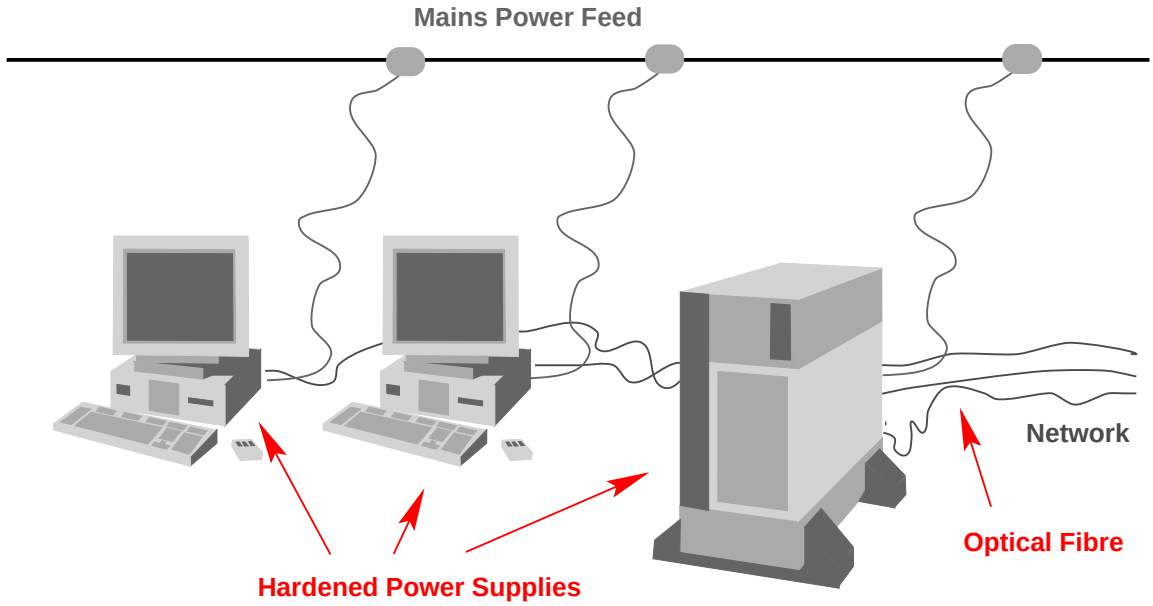

Fig.3.1 I/O and Power Interface Hardening

The more likely route for such equipment to appear in the near term will be through specialist suppliers, such as the large number of US manufacturers (Codar, Radstone, Cyberchron, AP Labs, Interstate Electronics) who supply "Milspec" packaged "ruggedised" versions of commercial computers to military and

government customers. Such equipment uses commercial internals in a robust military spec chassis, built to withstand vibration, high temperatures, impact, dousing with salt water and various other indignities which a computer in military service may have to suffer.

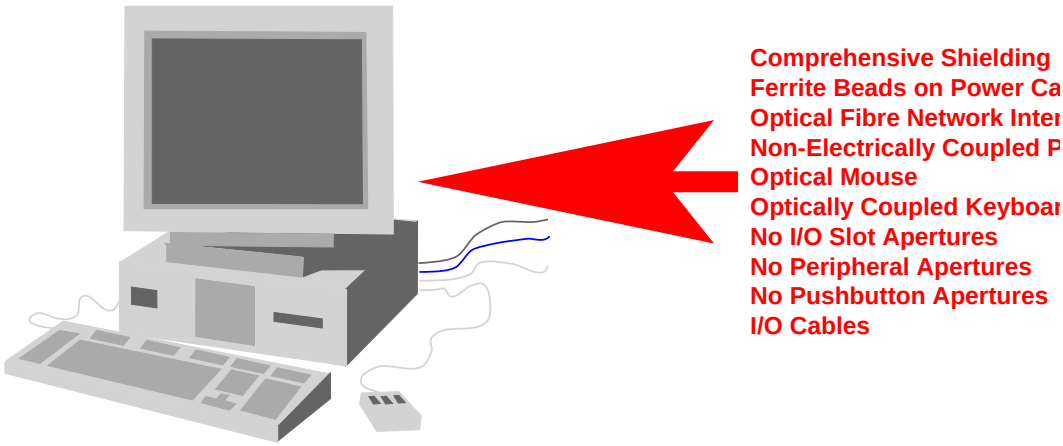

Fig.3.2 Comprehensive Host Hardening

An electromagnetically hardened commercial computer will not require the physical and environmental ruggedness of a military system, but will require the ability to cope with high RF fields, voltages and spiked power lines. Therefore, it is likely to be much cheaper than existing ruggedised equipment, but probably still much more expensive than off-the-shelf commercial equipment.

The first important area which must be addressed is that of providing high performance and comprehensive shielding for the equipment. Existing shielding for government TEMPEST rated equipment would probably be in the right class of performance. This would address the ability to keep to nasty RF fields from getting in through gaps, cracks and cooling grilles in equipment. Monitors, be they CRT or LCD based, will need conductive materials embedded in the screen or screen cover to render it opaque to RF radiation.

The power supply is the next item which needs to be addressed. A good measure of protection could be provided by a well designed conventional switchmode with high performance RF and spike traps designed into the hot end, but higher performance alternatives do exist. Many techniques can be used for non-electrical coupling between the mains and the low-voltage side of the equipment. The simplest is a miniature motor-generator scheme, where the "hot" mains side uses a squirrel cage electrical motor to drive via a shaft an internal alternator or DC generator, which is integrated with a regulator arrangement to produce the required +5V, +3.3V, +12V and -12V rails. Such a supply if built properly could easily fit into the form factor of most current tower case PC supplies. Equipping such a device with an internal flywheel would also provide a good resilience to short mains voltage dips. A power supply built this way would provide much smoother power than existing switchmodes, and also avoid the production of nasty RF interference, common to cheaper switchmodes.

More esoteric schemes may be used, such as hydraulic power transfer (mains driven pump to impeller to alternator), or fuel cell based schemes. In theory, such devices could be built with similar efficiencies to existing switchmodes.

Having addressed shielding, power and networking, we are left with the machine’s immediate interfaces, such as the keyboard, mouse, serial ports, external drives etc. In the short term, these may be problematic. There are no technical reasons why keyboards and mice cannot be built with optical fibre interfaces rather than copper interfaces. A keyboard could, in theory be built wholly with optical switching technology. Certainly the push toward fibre based SCSI interfaces, driven by demand for bandwidth, will solve the external peripheral problem, assuming that power supplies and chassis for external devices are suitably robust.

In terms of other interfaces, the characteristic backpanel’s worth of connectors will have to go, as these are all potential entry points. Again, there are no technical reasons why optical fibres cannot be used

for all connections between pieces of equipment. It would certainly solve many other problems our support engineers have to grapple with, such as ground looping between equipment.

The author has designed many computer chassis, and quite a few board level products (eg workstation motherboards, I/O boards, fibre comms equipment etc) over the years, and can state without fear of contradiction that the interface and packaging design changes required to cope with an electromagnetically hostile environment could be readily integrated into the existing technology base. Indeed, the commercial opportunities for smaller manufacturers in the production of hardened equipment chassis and interfaces, using standard commercial internals, are considerable in the medium term. We can hope that our industry will rise to this challenge.