Safety implications of industrial uses of internet technology

Prepared by
Tessella Support Services plc
for the Health and Safety Executive

CONTRACT RESEARCH REPORT
408/2002

HSE
Health & Safety Executive

Safety implications of industrial uses of internet technology

Dr. Mark English
Tessella Support Services plc
Robert Gordon House
Cavendish Avenue
Birchwood Park
Warrington
WA3 6FT
United Kingdom

This report examines the safety and health issues associated with the use of the internet and internet related technologies in and by industry.

There is a great deal of pressure for an organisation to take up the internet, outside pressure from the customers and pressure from within. The technology to allow connection of manufacturing and control systems to the internet is easily available, but there is no real evidence that adequate consideration has been given to the security and stability of these devices and systems. There are also disadvantages in exposing internal systems (business or control) to the outside world.

The quality of publicised safety-related data varies, and is potentially dangerous. Network integrity is paramount in any internet-based system that is used for control and automation, where a breach could have serious safety repercussions.

This report and the work it describes were funded by the Health and Safety Executive. Its contents, including any opinions and/or conclusions expressed, are those of the author and do not necessarily reflect HSE policy

HSE BOOKS

First published 2002

ISBN 0 7176 2268 1

Foreword by HSE

Internet technology gives cheap and flexible capability that can change existing industry practice and make possible novel practice, and has the potential to give rise to safety concerns. HSE commissioned Tessella Support Services plc to outline the safety implications raised by the industrial use of internet and related technology. The emphasis was firmly on current actual use of the technology. A degree of future-gazing was also included to cover credible foreseeable uses.

After detailed consideration of Tessella's findings, HSE has concluded that Internet technology raises no fundamentally new safety issues. However, the enhanced functionality and reduced costs mean that complex implementations are now more practical, and safety and security measures need to be thoroughly addressed. This would involve, in the context of safety issues, analysis of system, hardware and software and human factor issues.

Key to HSE's technical strategy is the international standard IEC 61508 "Functional safety of electrical/electronic/programmable electronic safety-related systems". Internet-related technical matters can be handled within this strategy, which offers a systematic approach to the safety of complex systems.

Through no fault of the contractor, the publication of this report has been considerably delayed. Some of the specific references may have lost some of their currency, but they remain valid illustrations of the author's arguments. To preserve confidentiality, the text refers to some sources of information as "site 1" to "site 6".

All opinions, statements and observations expressed herein are based on information from sources which the author believes to be accurate. While care has been taken in the preparation of this report, the author issues no warranty, expressed or implied, as to the accuracy or completeness of any opinion, statement, analysis or observation provided.

Edward Fergus

Technology Division, Health & Safety Executive

CONTENTS

EXECUTIVE SUMMARY

This report examines the broad safety and health issues associated with the use of the Internet and internet related technologies in and by industry. The report highlights usage of the Internet and related technologies across industry, and the safety implications.

Data transmission

The data transmission methods used in the Internet were devised to transmit mainly text messages in relatively low integrity applications. Traditional office applications have not required any fundamental improvement. However, newer Internet applications use data transmission in distributed tasks that require higher safety integrity. Safety applications include: remote monitoring, diagnosis, control and maintenance; full integration of computer systems covering both production and administration (thus introducing new vulnerabilities); publishing of safety-related information.

Remote access

In addition to the more traditional variety of computer networking, web-enabled controller devices are increasingly available, giving direct Internet access to equipment that would previously have been isolated. Although HS(G)87 identified the safety issues, the cost and inflexibility of the technology at that time discouraged the practice. Industrial uptake today is much easier. For an organisation that now has expertise in browser-based internetworking, this is an attractive approach to developing a complex industrial control network: the kit is readily available, cheap, and is not limited by a supplier's proprietary standards.

Complexity of systems

Encouraged by technology manufacturers, internetworking technology is readily available and complex systems can be assembled without fully understanding their vulnerabilities. Networks are also inclined to grow incrementally with new functions. Observation of actual applications suggests that inadequate "systems" planning is often devoted to network security and operating stability.

There is perceived to be an efficiency gain in integrating all management and production activities into a single integrated network – the so-called "shop floor to top floor" integration. This both adds complexity and introduces vulnerabilities. The growth of e-commerce will encourage this trend.

Security

System and network security has safety implications: data corruption can lead to unsafe decisions, and unauthorised control activities can be unsafe. Threats (viruses, malicious modification, and other breaches) are well publicised, but the threat and ease of malicious incursion is not fully comprehended. An organisation that integrates its management and production-control networks in a "mixed infrastructure" may be particularly vulnerable.

Domestic internet access will increase. Although web-enabled domestic equipment is currently an uncommon novelty, it is likely to become much more common as non-trivial applications are developed. Domestic applications will share the security vulnerabilities of industrial installations, but are less likely to be expertly managed. There is a possibility that domestic systems could be exploited to attack safety-related industrial systems.

Good network security working practices are known, and need to be encouraged in safety-related systems.

Network security is not a simple attribute than can be easily measured. A benchmark of practical assessment criteria exists that could provide the technical basis of a certification scheme to promote common standards.

Data gathering

When using remotely gathered data, care is needed to avoid making a bad safety decision from a wealth of irrelevant or unstructured data. Data gathering tools may also impose a pre-determined view that omits unanticipated failure modes, resulting in incomplete data. Ergonomics of data presentation is not trivial.

Publishing safety-related information

The legal liability for published information is unclear. Court decisions are few and inconsistent, and the status of safety-related data is untested.

Publishing safety-related information via the Internet raises issues of currency of information, timeliness, accuracy, applicability, editorial policy. There are technical factors that give rise to novel difficulties, and cultural attitudes that differ from conventional publishing.

Call centres and competency

Increasing use of call centres for technical support, remote monitoring of alarms and plant, fault reporting, and the clustering of safety-related data. Critical diagnosis skills reside in the call centre, while the local operator can describe the symptoms and apply the central recommendations. The competency of both the call centre operator and of the local operator is an issue. The central expert may misdiagnose through lack of a detailed local context and unanticipated failure modes.

Health related issues

A networked enterprise may develop into a distributed one, with greatly reduced co-location of collaborating workers and teams, raising issues of management control over the whole. A team of teleworkers may be “virtual”, and virtual teams in different organisations may collaborate. This approach may raise issues of work-related stress, isolation, and management supervision of safe working practices. Global and cross-border working raises issues of cultural differences, with the possibility for safety-related misunderstandings, and concerns over the fitness for purpose of goods and services.

1. INTRODUCTION

1.1 PURPOSE

This report examines the broad safety and health issues associated with the use of the Internet and internet related technologies in and by industry. The report will highlight the breakdown that can be made of the usage of the Internet and related technologies across industry, and how each has a different set of implications. Also to be discussed are the ramifications on management and business processes of implementing distributed systems usage.

The areas for consideration will be broken down into smaller groups. For consideration then is the use of the Internet, as commonly understood, as a part of an information or control infrastructure, and the parts of the internet technologies in use within organisations. The effect of implementing the technology will be considered from the perspective of how management style can change, and how ground floor working practices can also change.

The report also briefly looks into the way internet technology is being gradually introduced to domestic life, and possible implications of its usage.

1.2 SCOPE

The broad scope of the report is the assessment of Health and Safety related issues to do with the use and uptake of Internet related technologies in industry. The effect of the distribution of information across an organisation (or co-operating organisations) on operational procedure, and any decisions based upon that information is also to be discussed. Briefly considered are the options for domestic usage of internet technology in the arena of control and automation.

This report will not cover low level technologies such as best practice for developing communications protocols (for example), or the details of remote operation as these are covered in a separate reports [Lafave (2000), HSE (1995)].

1.3 CONTEXT

This report forms part of a general inspection of IT and communications in industry being undertaken by the HSE. This report will aim to identify potential areas where the uptake of internet related technology for information distribution (such as the provision of misleading, or faulty data), or the corporate integration of data and manufacturing processes (with mixed usage of standard infrastructures) could cause a problem. During the course of preparation of this report various methods for guarding against some of these problem areas were identified, and these will be mentioned.

2. THE INTERNET

2.1 WHAT IS IT?

2.1.1 The Internet

The Internet has its origins in a military requirement for a communications network that is resilient to a degree of damage. The Internet is a complex system of computers interconnected by a mix of dedicated land and submarine cables that provide multiple transmission paths from message source to destination. The communications network consists of hardware (cabling and routers), and software (the protocol that is used to manage the information flow on the network). The protocol that is used for the Internet is called "IP" - Internet Protocol - and is usually used in tandem with "TCP" - Transmission Control Protocol.

The TCP/IP protocols are of the "packet switching" type. At the source a message is cut up into fragments ("packets") that are numbered in sequence, and are marked with the destination and with any special transmission requirements. The complete collection of packets is routed through the network to the destination, where the original message is reconstructed. The TCP/IP protocols protect the data against scrambling during transmission and reconstruct the message at the destination.

In addition to sending data over long distances, it is frequently required to connect geographically close devices and workstations into a "local area network" (LAN). An "ethernet" is a very common technique for local area networking. Strictly, ethernet refers to a particular type and definition of data transmission protocol, but this report uses the term colloquially to indicate a variety of local area network protocols and transmission media. Many industrial controllers and workstations are equipped with an ethernet interface as standard, allowing them to participate in a distributed system.

Communication by Internet TCP/IP and ethernet in local networks is non-deterministic. You can specify where you want your message to go, but you can't say when it should arrive, or by what specific route it should get there (in general). Contrast this with a LAN where the path of the message is governed by the network topology. The Internet achieves the required resilience in the communications network by routing the message by whatever path is available. If a message cannot get through by a particular route because part of the network is broken, then the message is re-routed through another part, and this will continue until the message gets through.

Other protocols worth mentioning in this context are "RealTime protocol" (RTP) [Schulzrinne *et al.* (1996)] and "Resource Reservation Protocol" (RSVP) [Barden *et al.* (1997)]. These protocols find application in the voice, audio and video transmission areas of internet use. RTP provides end-to-end network transport functions suitable for applications transmitting real-time data such as audio and video, whereas RSVP allows an application to reserve enough "space" (called bandwidth) on the network. Both of these give some level of guarantee of prompt network packet delivery. Although the use of TCP/IP currently by far outweighs the use of these two protocols, they are likely to become more prevalent as internet technology is taken up more widely.

2.1.2 The World Wide Web

Every computer in the Internet has a unique address in a format that is dictated by the TCP/IP protocols. Where a computer contains a collection of documents, a further addressing scheme is needed to identify these items uniquely in the Internet. Such an address is a Universal Resource Locator (URL). The World Wide Web (WWW) is essentially an addressing scheme that allows the entire Internet content to be accessed as a collection of cross-referenced documents, each

uniquely identified by a URL that hides the details of the TCP/IP address. In short, a URL is a unique address in the WWW.

In addition to its own text or other content, a document may contain URL references ("links") to other documents in the WWW. A link can point to a document anywhere in the Internet. In this way, the content of several documents at different physical Internet locations are logically combined into a single WWW document.

A user needs a web-browser (such as Netscape Navigator, or MS Internet Explorer) to read a WWW document. The browser displays the text or other content of the document, and is aware of any URLs that are embedded in the document. When the user "clicks on" or otherwise activates the embedded URL, the browser makes an Internet connection to fetch the content identified by that URL. Thus the user has access, through the web-browser, to any information source that has been published on the World Wide Web.

2.2 NETWORKING BASICS

Since the early days of being a network for researchers, the Internet has taken on many roles. It is currently used for email, information distribution, archiving, video conferencing, direct commerce and much more. This section introduces some terminology.

2.2.1 Local Area Networks

Within a company that uses networked computers there will be a Local Area Network (LAN). This LAN is more than likely to use ethernet technology. This is a well known and consequently well understood and trusted physical network technology. In order to carry information from computer to computer a protocol is used on the network to ensure data integrity and no data loss. The detail of protocols is covered in [Lafave 2000] so will not be considered further.

The protocol generally in use in office and corporate environments is TCP/IP. There are other networking protocols, but TCP/IP is by far the most prevalent, and is under consideration by many manufacturers of automation hardware as suitable protocol for automation and control [Reeve (2000)].

The migration of industrial protocols onto ethernet and TCP/IP continues [Lock (2000)], with Profibus and ControlNet for example declaring their aims to develop TCP/IP capability. The apparent attraction of TCP/IP is the customer "push" towards standardisation, whereas suppliers would prefer more proprietary networks [Lock (2000)]. The commercially available components for making ethernet connections are suitable for industrial automation [Reeve (2000)]. However the devices (hubs, switches) have to be made to withstand the industrial environment, which is harsher than the standard office environment.

2.2.2 Use of an Internet Service Provider

Connection to the Internet is through an organisation called an Internet Service Provider (ISP). This service acts as a portal for information to and from the Internet. Connection to the ISP can be through a leased line, ISDN, standard phone line (a so called dial-up connection), or through one of the new media such as ADSL (Asynchronous Digital Subscriber Line).

2.2.3 Private Networks

A Private Network (PN) is intended to restrict communication to a defined group of users. The PN may make use of some publicly available communication facilities, but it will do so in a way that restricts access to the chosen group.

This report does not cover the specifics of setting up a PN. However, the following examples are illustrative.

- A typical PN technology is the use of leased lines between the geographically distributed offices of an organisation. These are permanent connections over the public telephone system. This has some implication for security, since the connection goes through public exchanges, and so can be tapped.
- Another method of forming a PN is through the use of line-of-sight microwave connections.
- In contrast, ISDN connects on request for service; one machine will dial up another to make the connection. This has some implication for the availability of the service. Potentially, a different organisation could dial into an ISDN phone number when that phone number is not in use, thus making it unavailable.

A Virtual Private Network (VPN) offers greater privacy. An organisation with geographically distributed offices can connect them together over the public telephone system (e.g. using ISDN), and then additionally use strong encryption to ensure that all data transmitted between nodes (branch offices) remains private.

The formation of a VPN through the use of encryption can also take place over the very public Internet. The privacy results from the permitted nodes being the only parties able to decrypt the data, and the "virtual" status is because this solution uses a physical public network for transport.

To make use of the Internet from a private network of any sort, a connection to an Internet Service Provider (ISP) is required. The ISP link always raises the possibility that the private network will suffer unwanted incursion, because the ISP link provides a potential entry for any unauthorised but technically competent parties.

2.2.4 General Internet services from the ISP

There are many services which an organisation can use over the Internet. The common examples include email, FTP, gopher (primarily a service for indexing and searching for text documents), and WWW browsing and publishing.

A recent development is the emergence of voice and video communications over the Internet, and currently much work is being done to develop services of "voice over IP".

All these services can be used in a two-way fashion, with data flowing into and out of the organisation. Alternatively, the networking may be required solely to link the geographically distributed parts of an organisation, but with no outside access. This report considers these two illustrative models of an organisation:

- A "public" enterprise that is willing to exchange data, email, etc. to anyone who contacts it.
- A "private" enterprise that uses internet technology to facilitate communication across its distributed parts, but accepts no contact from outside its own organisation.

A practical organisation is likely to have both private and public aspects, but this report illustrates only the two extremes. The public option will consider the cases:

- An organisation using two way services - email, WWW data serving and browsing.
- An organisation using the WWW as a distribution medium, leaving data to be held by an ISP.

2.2.5 Private network access models – open and closed

From within a private network people generally do not have individual, direct access to the Internet from their desktops, but usually these services are routed from the individual desktop computers through gateways on the network. These gateways manage the connection(s) to the Internet.

This report considers the general situations of a private network being open to bi-directional information flow with the Internet – the Open model, and that of a closed network – the Closed model. These two network models are illustrated below.

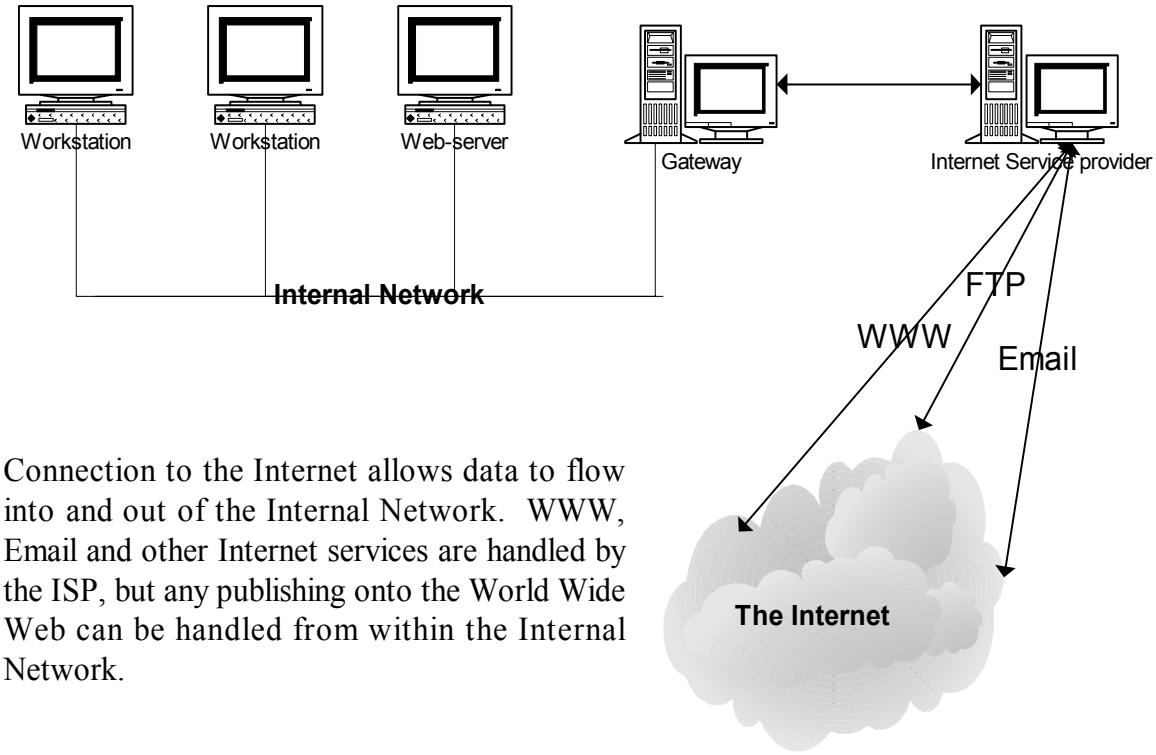

Connection to the Internet allows data to flow into and out of the Internal Network. WWW, Email and other Internet services are handled by the ISP, but any publishing onto the World Wide Web can be handled from within the Internal Network.

Figure 1. The Open Network model

In the Open mode, the internal network is open to bi-directional data exchange with anyone on the Internet. Examples of bi-directional data could be email, WWW browsing, and FTP.

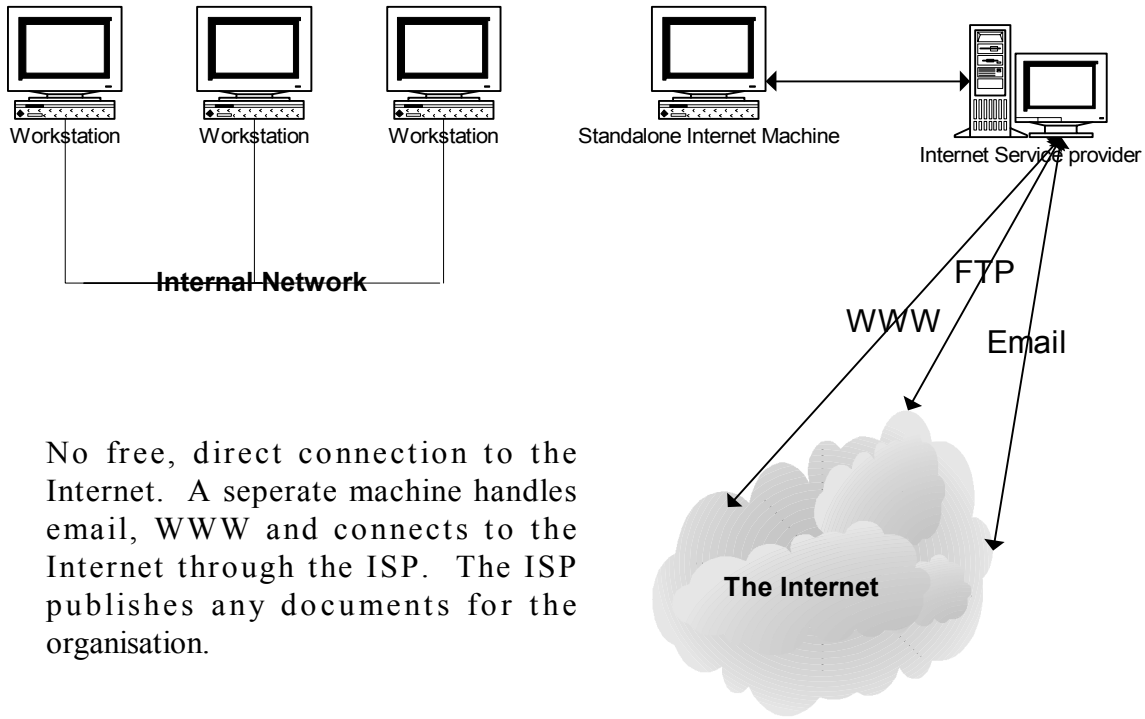

Figure 2. The Closed Network model

In the Closed model there is no direct, free connection to the Internet. All interaction with the Internet is performed through a single machine or dedicated network.

3. USING THE INTERNET

3.1 THE "PUBLIC" ORGANISATION - PART OF A PUBLIC ENTERPRISE

This section introduces the concept of organisations who are using the "Public Internet" as part of their daily business. For example, bodies that make available data and resources for access via common web browsers, and have an email gateway to allow outside email into the company. This describes any system that is open to interaction from anyone on the Internet.

Using the "Public Internet" as part of business includes the following areas:

- Publishing of data
- Web browsing
- Email

The publishing of data can be entirely hosted by an ISP (that is, the data being published is stored and distributed from a computer that is external to any company network), or can be from within a company network. Both the Open and Closed access models are appropriate:

- [Closed model] Web-browsing can be done from an individual machine connected to an ISP via a dial-up or ISDN connection, which separates this activity from the main business of the internal network.
- [Open model] There are potentially many points of contact from within the company network to the Internet.

Email can be dealt with in a similar way to web-browsing:

- [Closed model] This facility is available from a single machine only.
- [Open model] This facility is generally available many machines.

3.1.1 Data Security

Data Security should be thought of as susceptibility to loss or corruption through (1) directed action, or (2) indirect action that is the knock-on effect of a directed attack elsewhere, (3) through accident, or (4) inappropriate software controls.

3.1.1.1 Directed Action

The subject of security against "hackers" has been ongoing since operating systems and applications that used networks were invented [Cheswick (1994)]. A directed action could be an attempt to enter the company network to browse, modify or steal data.

The first category is by an individual, and is a one-to-one activity (one agent, one session). The statistics on successful break-ins to what are supposed protected systems are staggering. Randell *et al*. (1999) report that the US Department of Defense suffered more than 165,000 *successful* attacks in 1995. It should be noted though, that in order to be attacked, an organisation must be the electronic equivalent of visible: the Open access model is highly visible, while the Closed model is much less so. An organisation is also more susceptible to attack if it is well known and with a high profile. Unix is more prone to abuse in this fashion than Windows operating systems as it has been around a lot longer and has the in-built facility to allow remote use.

A second category of attack is the one-to-many (one agent, many instances). This sort of attack can include viruses and code carried as a "payload" in email or web-pages. A virus is a program

that replicates and distributes itself on your computer. It does not generally have to be aggressive - it may just replicate until your computer is so "infected" that the computer is unusable (e.g. hard disk full). Alternatively, an aggressive virus may have a destructive payload (i.e. a small utility that is carried by the self-replicating virus body) which may operate in one of many devious ways to destroy data in your system (e.g. reformatting your hard-drive).

A virus can be introduced through execution of non-validated software (e.g. screen-savers, games) downloaded from the Internet, or simply through the careless use of floppy disks. This method of gaining entry to a PC is also used by so-called "Trojan" programs. These sit on the computer and feed information back to a remote host. Some transmit keystrokes (hence gaining passwords), some send password files back. These Trojans can often form the first part of a concerted attack [Cheswick (1994)].

A new variant on these low-level viruses makes use of vulnerabilities in operating systems and embedded scripting languages in applications such as web-browsers, office tools, email clients and such. The recent "Lovebug" email virus was a malicious self-executing script that took advantage of a facility called "Active Scripting" - essentially an executable program that accompanies the email message - that performs an action and then proceeds to proliferate via email to other parts of the network. This type of attack is likely to become more frequent. The "Lovebug", the "Iexplore.zip" worm, and the "Melissa" incident are all well known attacks within the past year.

Windows operating systems are more prone to these last two attacks than Unix. Note that Unix is not necessarily inherently strong in this respect, but rather that Windows equipment is more common and therefore a more attractive target. This may change soon with the growing popularity of Linux systems.

Another attack that can affect data distribution is a "Denial of Service" (DOS) attack. A DOS attack (and a Distributed DOS attacks i.e. a many-to-one attack that comes from several sites concurrently) does not necessarily enter the network or computer hosting the service, but it prevents that service from effectively performing its task. One way to do this is to flood the system with bogus requests, thus congesting the system for legitimate requests.

A more devious Denial of Service will take advantage of some known vulnerability of the system. For example, it was recently identified that if Microsoft Information Server receives a badly formed request involving many suffixes (e.g. www.company.com/this.is.wrong.this.is.wrong.bat), it then spends much effort searching its database to recognise the suffixes. If many such badly formed requests are received, the Server is monopolised trying to resolve the requests, resulting in a denial of service to legitimate users.

Direct attacks can affect specific processes (e.g. the delivery of safety information), or the computers and operating systems they are running on.

3.1.1.2 Indirect attack

In an indirect attack, a service that is running on computer X can be brought to a halt because a related computer Y (with which X shares some resources) has come under attack.

For example, the recent Lovebug incident brought many corporate networks to a stand-still as the level of email traffic hit record limits, and swamped the network, thus locking out other networking applications. Clearly this stopped all inbound and outbound email, but any systems that were trying to share resources, such as networks and disk space would also have been affected.

3.1.1.3 System Modification

The system hosting a service related to data delivery can be modified in a variety of ways that will leave that system in an unknown state. Modifications can include (for Microsoft Windows) additions to system files, working files, program files and the registry (i.e. operational settings). The effect of modifications can range from no effect at all to a service becoming unusable or unavailable, with various states in between which may yield undefined behaviour.

Systems can be adversely modified through the careless upgrade of system software. Protection against accident of this kind is very difficult, and relies on having appropriate recovery procedures in place to ensure that any interrupted service can be restored rapidly. This is best addressed by adopting good practice in system management, and is not considered further in this report.

Systems can also be modified either through the hidden action of agent software, or by directed action. Of interest here are programming techniques that update, alter or reconfigure computer systems without wilful instigation of a legitimate user.

Technologies that can potentially modify computer systems are ActiveX, Active Scripting, and Java. The first two only affect Windows operating systems. Java will run on any operating system that runs a modern web-browser. In all cases, the essential problem is that an executable program that is downloaded from elsewhere in the network may run on the local computer without the knowledge or express permission of a legitimate local user.

ActiveX is a component-based technology (i.e. individual program designed for re-use) that was initially specified by Microsoft for use within web applications, and is based on the "Component Object Model" (COM) that describes the standard interface between Windows components. Consider a browser on computer A that is accessing a web page somewhere in the network. If the web-page requires the use of a particular ActiveX component that is not available on computer A, then that component is downloaded from a web-site on another computer B and is executed on computer A. Depending on security policies implemented in computer A, this component can be downloaded and executed without the person viewing the page ever knowing. This ActiveX component can then execute with all the privileges that apply to the local user: it can access the hard disk, files, send email, move local data to or from another site, or it could deliver a virus. Although there is some standard protection - Microsoft provides a certification mechanism to filter components from trusted locations only. However, users have to be sufficiently knowledgeable to make informed use of it.

Active Scripting is the technology behind the Lovebug email incident. Active Scripting can access your files, change registry settings, write to hard disk and cause emails to be sent in your name. There are also ways of manipulating, or writing WWW pages with Script embedded in them [Moody (2000)] that when viewed can cause your office applications to perform unwanted operations, or can cause emails to be sent pertaining to be from the user of the web-browser. Again the extent of damage by this technology abuse can be limited by informed use of security settings.

There are outstanding problem reports from Microsoft on these problems which can be viewed at http://www.microsoft.com/technet/security (Accessed 18/05/2000).

Java has generally caused less problems than ActiveX and scripting mainly because it was designed to be secure, to run in a protected area of memory (called a "Sandbox"), with limited access to the host computer's resources. Java does not allow access to the local files for example, and a downloaded Java program (an "applet") running in a web-browser can send data back only to the server from whence it came. If functionality is required that breaks these simple rules, then the user has to make explicit changes to the web-browser. However, as Java

use becomes more commonplace, and as greater functionality is demanded, the likelihood of malicious use increases.

The effect of these system modifying technologies may be to affect a safety-related system (1) directly if it is hosted on the same computer, or (2) indirectly through denial of an important shared resource (such as network or storage).

3.1.2 Data Availability

In this section it is assumed that the system that is the source of safety-related data is itself stable, and is not liable to attack. In this context, data availability covers responsibility for the data and its delivery, the level of the data given, its pertinence, timeliness, and that it is not misleading.

3.1.2.1 Legal aspects

Safety-related information delivered over the Internet is transported through an ISP. Whether the model being considered is a Closed network (all data hosted at ISP) or an Open network (data hosted at company connected to ISP) there may be a question of shared responsibility. Who is responsible for ensuring that the safety-related data is available when needed: the company that controls the content and is effectively broadcasting it, or the ISP carrying it?

There has been a small number of court cases concerning responsibility for data hosted at ISPs, but the position remains unclear.

- In a French case [Reuters (2000)] an ISP was inadvertently hosting objectionable extreme political material. It was decided that the ISP had taken reasonable measures to monitor the content, and that although some material had escaped attention, further technical measures were unnecessary. The French ISP was held not liable.
- In a German case [Gray (2000)] it was decided that an ISP is liable for facilitating the illegal copying of copyright material. The argument appears to be that the German ISP is more than a transparent re-broadcaster of information like the phone system, and must take some responsibility for the content.
- In a second French case [Reuters (2000)] an ISP was ordered to block access in France to material that is illegal in France but was brought from US web sites where it is legal. The French ISP appears to be responsible for the content.

Although none of the above cases directly concerns safety-related information, they highlight the issue of who is responsible for the data content - the person or organisation putting the material on the servers, or the ISP that hosts and distributes it. Note also that the variety of rulings means that any data distributed via the Internet may have strong legal geographic concerns.

The legal status of Internet and WWW data is unclear, in as much as it crosses international boundaries, and it appears to be undecided as to who is responsible for what. Consider the scenario of a French company that uses an American ISP to host their web site, and that web site is accessed through a UK ISP. A UK citizen could make a safety-related decision based upon inapplicable information supplied via the WWW from this site. If unsafe consequences result, who would be liable - the French company for making the information public, the American ISP for publishing it, or the UK ISP for delivering it to the user?

3.1.2.2 Currency of data

Where an organisation makes a safety-related decision based on web-based document or data delivery, then it would clearly have to consider what to do in the case of a power-cut (or one of the many other forms of loss of service) that would make the remotely hosted data unavailable. Some of these interruptions will be beyond the control of the user organisation.

If the data is critical to the user, should the organisation that provides the data consider a backup service in the case of their main service failing?

3.1.2.3 Timeliness

The timeliness and correctness of data can have two influencing factors: caching, and editorial policy.

Caching is the practice of storing web-pages locally for faster access, as opposed to fetching them from the definitive Internet source on every access. Pages can be cached at the client (i.e. the web-browser) side, or at the ISP.

When using a web browser that has cached pages for faster access, it is possible to be reading out of date safety-related information. This can be alleviated by informed use of the browser (typically, forcing it to reload pages automatically from the web site).

An ISP may cache web-pages to give a faster perceived service to the clients. It is often observed in practice that this results in old pages being available by accident. This caching is beyond the control of the ISP user, who must discuss the problem directly with their ISP to arrive at a solution.

If an organisation manages its own web site and has control over the browsers accessing the site, then it can implement schemes to ensure the timeliness of the data being read. A typical solution would be to institute a policy of refreshing every 5 minutes all the pages cached on the web server.

A related problem concerns the varying capabilities of the web browsers to display all the information on the page. This is becoming more pertinent as browsers have diverging functionality, for example Microsoft Internet Explorer runs VBscript, whereas Netscape Navigator does not. Thus, any pages using VBScript to format, modify or import data will not run on a Netscape browser, and a user could read a different set of information depending on what browser is used to view it. The safest option is to use pages constructed with the lowest common denominator in terms of content formatting and control. A prudent policy is to audit all the machines to ensure they can all view the same pages and data.

An obvious but sometimes overlooked requirement is that the supplier of the data should have an editorial policy that keeps their data up to date, and pertinent to their product (for example). As a practical illustration, several sites relating to potentially hazardous products were inspected by way of example. The quality of the data delivered varied considerably. The sites 1 and 2 dealing with resins and chlorine-based products were found to have good references to regulatory standards, and supplied the data that could be assessed against a defined standard. In contrast, site 3 had not updated its data about its epoxy product since 1993.

3.1.2.4 Accuracy

Site 4 had references to HSE guidelines, but also had a typo in naming one of the volatile organic compounds. From a safety perspective the naming of hazardous materials should be accurate.

Another example of potentially inaccurate data on a web-site (site 3) is the inclusion of a disclaimer to allow them to vary the product. This combination of out of date information and also the fact the data may not pertain to the product obtained renders it not only useless, but potentially misleading. While site 3 contained references to datasheets for the chemical components, it was not clear how to link to or obtain the datasheets.

The referencing to standards for data supplied could be particularly important in order to make a health or safety-related decision. Site 5 (a hospital) had a recommended schedule for a drug dosage, but without making clear the source and authority of the recommendation, nor the intended recipients of the recommendation. If it was for use by medical staff, it should have (1) a date, (2) who recommended the dosage, and (3) some other audit trail information. If it was not for use by medical staff then why was it available? This is potentially misleading.

Another aspect of accuracy is the choice of language. The Internet allows distribution of safety-related data world-wide. Clearly this data must be understood comprehensively, and so data must be supplied in a language that is acceptable in the region where the data is accessed and used. Some sites have versions in several languages, so the user can choose which language to browse in.

Geographic applicability of safety-related data may also become more important as organisations trade globally and distribute information globally. Where data has some regulatory or legal significance (e.g. Occupational Exposure Limits under the UK COSHH regulations), and where different regions may set different acceptable variations or representations, then it should be made clear to the user where any data supplied via the WWW is pertinent.

Examples include:

- [NASA(2000)] The failure of the Mars Climate Orbiter was attributed in large to a failure to convert measurement units from one geographic standard to another. This highlights the importance of cross-geographic definitions, and how distributed systems development needs as much review and engineering as a local one.
- [EPA (2000)] The US Environmental Protection Agency produces factsheets on air toxicity regulations but does not state where they are applicable. Though this factsheet can be accessed from the EPA homepage, a search of the Internet would bring up the single page of facts, which has no clear indication of where the data is applicable. Someone charged with locating regulatory information or guidance could then make use of inapplicable data from the Internet in a safety-related decision.

3.1.2.5 Editorial responsibility

As web-pages are linked through “hyper-links”, whereby it is possible to click on a reference to open a remote page, the question can then arise as to who is responsible for that link. For example, if site A’s page of safety data has a link to site B with further pertinent information, is site A or site B responsible for ensuring that the link points to a currently available page? An example of this problem was found on the HSE site [HSE (2000)], where the link to the Corgi gas site was a “dead link”.

For a complementary viewpoint, consider site A that links to another site B that has a health and safety resource. There is no practical way for site B to know about that reference, or that any other site is depending on it. If site B then issues new versions of its standards or safety-related information, but does not clear out the old ones immediately, then it is possible for site A to receive incorrect data.

3.1.2.6 Internet publishing compared to paper publishing

It seems plausible that a major impact would be made on the above problems by applying to electronic media similar standards of document control as are applied to paper. The procedure for reviewing Internet publishing content should be as rigorous as that for paper publishing. It is often observed that there is a general informality in dealing with the Internet, including the reviewing of material prior to being published on the Internet. Though people may review documents, there is the attitude that as the media of publication is electronic that "things can always be changed later". In the case of publishing safety-related data, later may be too late.

3.1.3 Technology uptake

The use of the WWW to distribute information is widespread. Many companies use it to distribute product information, and it would seem from the selection of sites inspected that the majority are cautious, they use informative disclaimers, and they reference the relevant regulatory bodies and standards.

There is no one common approach observed to formatting of data, audit trails, references, though some like Site 6 did appear to have copied their paper datasheet exactly judging by the format, so appearing more comprehensive.

From discussions with staff at Site 0 about their chemical emergency procedures it became clear that they had a system that it would currently be very inappropriate to implement as an Internet based system. The procedures for a chemical emergency (there are several chemical works in the area of responsibility of the Site 0 organisation) are encapsulated in a book that is copied and distributed to several emergency services and organisations involved in clearing up spillages. This book is signed out to the organisations so that the Site 0 organisation can control who has authorised access, and who is on the update list. There are several reasons why this set of procedures could not be moved to an Internet system. Firstly, not all the organisations have Internet access. Secondly, the people who work with the procedures may need to carry them around with them, or take them home to read. The staff that have to implement the procedures apply a strict interpretation and they like to have "The Book" on hand. If this system were moved to an Internet based publishing solution then there would be problems: with system usage, people printing off copies, ensuring that the designated people had received the latest updates, as well as some people not being able to access it.

From inspection of the WWW sites available it is not clear how many are based on an Open network solution, and how many are based on a Closed network for their web publishing services.

3.2 USING INTERNET TECHNOLOGIES WITHIN AN ENTERPRISE

This section will discuss an organisation that runs a private network and uses the Open network model. The discussion will cover the use of internet technology within the organisation, and the influence of bi-directional data exchange with the Internet on that network.

3.2.1 The internal network

The general situation within any networked office environment is along the lines of independent co-operating machines. Each machine is independent, has a standard set of office applications, and is networked with file and print sharing facilities.

3.2.1.1 The Office background

These machines are generally installed with email tools, WWW browsers and usually a suite of office productivity tools (such as word processor etc). Depending on the choice of Operating

System (e.g. Windows, Unix), and the supplier of the 'desktop tools' (e.g. Microsoft, Lotus) there will be different vulnerabilities to the various attacks described in earlier sections.

3.2.1.2 Use of Freeware and Shareware

Within internal networks a company may make use of downloaded freeware or shareware in order to minimise costs. These softwares being of unknown pedigree (SOUPs) would have to be validated against a strict set of tests to make sure that they were fit for purpose. This report is not going to explore software validation, but it should be highlighted that being connected to the internet makes it much easier and tempting to obtain and to use shareware and freeware. Clearly, using software that has not been validated in a system associated with safety data (or systems), is introducing uncertainty. Using SOUPs in a system that shares resources (e.g. network, memory, disks) with a safety system could also have a knock-on effect on linked systems and thus introduce the possibility of linked failures.

3.2.1.3 Internal Information Distribution

The considerations that apply to the distribution of data via the Internet also apply to internal information systems (i.e. an intranet) for the dissemination of safety, quality or procedural information. Some companies no longer have paper versions of their QMS and emergency procedures. What happens in the case of a power-cut or similar common-cause failure that removes much or all of the network, or even the scheduled maintenance of a server? How can people access the relevant data in order to make informed decisions in an emergency?

With greater reliance on networked servers, and shared resources such as hard-drives, or work-flow/groupware software in the move to a "paperless office", the process of protecting the local data from corruption cannot be emphasised enough. For example, if the power fails when you are completing a safety audit using a paper based system, you have not lost your work. However, with an intranet/network server based system, if the power drops or if the hard disk breaks down, then potentially all the data stored centrally is lost. There are organisations who offer a special technical service to recover data from damaged disks, but this is not a recommended approach to data protection; an effective policy on data back-up is much preferred. An effective procedure for a regular back-up will also protect against the process of accidental data corruption (which would ruin any sort of audit trail), as it is possible to recover data from the previous back-up.

3.2.2 Manufacturing and Automation

As manufacturing has progressed over the years it has relied on larger amounts of automation to reduce operating overheads. There is financial and commercial pressure to adopt internet technologies into their automation and control processes, giving rise to vulnerabilities that potentially affect safety.

3.2.2.1 Mixed Infrastructure Usage

A current trend in industrial hardware control is based on data networking using semi-open standards such as Fieldbus, which has many variants [Babb (2000a)]. A more detailed study on available buses is Lafave (2000). A related trend is the increasing use of ethernet and TCP/IP as a connecting medium for discrete, and previously separate, subsystems [Reeve (2000)]. The reasons for this will be discussed later in this report.

In the so-called "shop floor to top floor" approach, the whole organisation is integrated into a single network. Where an existing corporate network (that already supports office applications) is linked with automation and control systems (such as Fieldbus and TCP/IP), experience indicates that safety may be compromised [Mintchell (2000)]. The load on the network from

office work may interfere with the need of the automation and control systems for timely responses. It is suggested [Mintchell (2000)] that the network should be split, and smart routers used to separate industrial command packets from the more mundane communications.

Another aspect of mixing the infrastructure usage is susceptibility to downgraded performance of the automation processes as a result of sharing resources with the office systems that have certain vulnerabilities. For example, if the office systems pick up a virus that floods the network with email, thereby monopolising bandwidth, then any other services using the network will be downgraded.

The design and implementation of a network is not a one-off activity. It is quite usual for networks to grow "organically" in terms of hardware and data traffic. What might begin as a satisfactory mixed infrastructure could be degraded over a period of time as new network devices and software are added to a system. Performance problems with critical systems may appear unexpectedly in a previously problem-free installation. To avoid such problems, frequent monitoring and evaluation of a mixed infrastructure is necessary, together with strict controls on the addition of any new hardware or software to any part of the network.

3.2.2.2 Buses and Embedded Devices

When controllers and devices (PLCs, motors, servos, etc.) are networked together (either directly, or via a gateway to an ethernet local area network), can they be affected by viruses and similar sources of interference? In order to affect a networked device, the virus must be capable of executing in an operating system on the device. Currently most devices are unaffected as they have proprietary operating systems, or are very low level devices with little scope for remote modification.

Embedded controllers may be more susceptible if they depend on higher level operating systems such as Microsoft WindowsCE (soon to be repackaged as PocketPC). The underlying organisation of WindowsCE is very similar to the well known MS Windows operating systems, and so could potentially display the same susceptibilities. This will depend on the actual configuration of the embedded WindowsCE system in the controller, what application programs are available on the controller to be exploited by an attacking program, and what degree of interaction the embedded controller has with web-based technology.

Knowledge of WindowsCE, and the use of industry buses for communication is limited, whereas knowledge of the Internet, Unix, Windows 95/98 and NT, and TCP/IP is widespread. This reduces the likelihood of any directed, or one-to-many attacks destined on embedded CE based devices, or on any other device that is connected via an industrial bus to a local network. However, if the uptake of industrial network connectivity increases then clearly the knowledge base will increase, and so will the likelihood of malicious or inadvertent attack. In short, non-Windows systems are not necessarily inherently safer – they've just been subjected to less malicious attention to date.

The network configuration chosen for connecting the industrial network to ethernet will also have a bearing on the potential for interference. Consider a management application that resides remotely on a TCP/IP network and queries the status of an embedded controller. Figure 3 shows an embedded device that is locally networked into an industrial Fieldbus, which is itself linked to the wider network via a gateway. The gateway forwards the query from the management application to the embedded device. In contrast, Figure 4 shows an embedded device that is connected directly to an ethernet, has its own TCP/IP address, and receives the query directly.

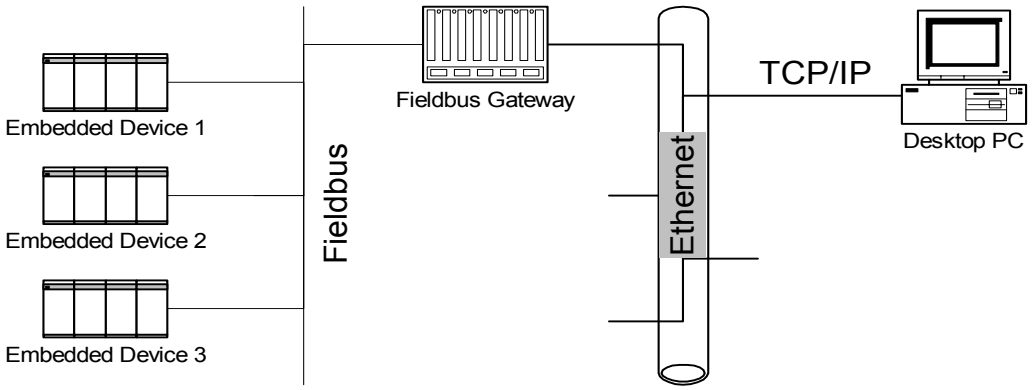

Figure 3. Access to embedded devices through a Fieldbus gateway from a Local Area Network.

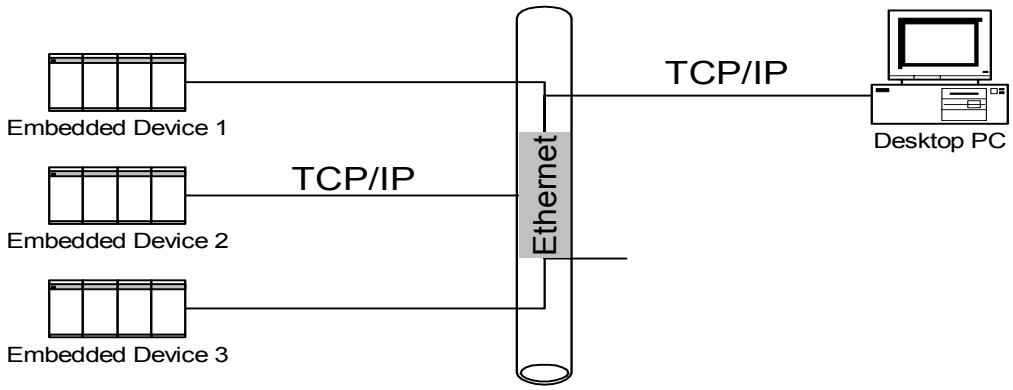

Figure 4. Some embedded devices allow direct connection to a LAN ethernet running TCP/IP.

In Figure 3 any vulnerabilities of the embedded operating system would be difficult to exploit unless the gateway vulnerabilities are also well enough known to be susceptible to attack. In contrast, in Figure 4 all of the standard hacking tools could be used in order to try and enter or break the embedded device.

The vulnerability of the Windows-based applications to Active Scripting, or running of ActiveX components, can be generalised to a problem of running unvalidated software. A technical/management policy that configures embedded controllers to run only authorised software would greatly reduce this vulnerability.

In a mixed infrastructure of administration applications and embedded controllers, unvalidated software can be introduced through a channel that is itself unaffected by any harmful features of that software. For example, a Unix machine acting as an internet news-server could fetch a news article that contains text that is interpreted by news-readers (such as Outlook Express) as Active Script – essentially an executable program that is downloaded with the message text. Unix machines and associated news-readers would not be affected by these Scripts, but PCs and PC-compatible controllers would. In this way, viruses can be introduced inadvertently through a different technology type.

3.2.2.3 Controllers using TCP/IP

A new generation of controllers is available that interface either directly to the Internet [Babb (2000b), Siemens(2000)], or that communicate directly using TCP/IP.

The aim of making controllers communicate directly as web-servers onto the Internet or Intranet is to permit direct remote access to data held in the controlled device. Enabling

technologies (such as the "Mimic" product from Dedicated Engines, or "Transparent Factory" from Schneider Electric) allow current installed systems to become "web enabled". The main difference between these two products is that Mimic is a one-way publishing of SCADA or control data onto the Intranet/Internet, whereas the 'Transparent Factory' is bi-directional so will also allow control. The "eMis" product from Dedicated Engines also allows the configuration and control of remote instrumentation. Remote operation is covered in HSE (1995). A variation introduced by this tool is the capability for user-specified safety features such as soft interlocks.

As these new devices are implementing a known internet technology, they can benefit from the experience the web developers have built up over the years. As these devices will be controlling processes as well as monitoring them, they could at least be proof against all the types of attack that have been noted historically. The CERT (Computer Emergency Response Team) website (www.cert.org) has recorded many of the vulnerabilities of webservers.

Some devices such as the "Procidia" system from Siemens is said to use a "hardened webserver" with "built in security" [George (2000)], but what does this really mean? A practical interpretation would be that in order to ensure a level of service for these network mounted controllers, good practice should assess them to a "proofing level", whereby they are certified proof against a judicious selection of the webserver problems known to CERT up to a specified date. This would be in effect a type of security benchmarking.

The data available from a web-enabled controller is liable to be misunderstood, or to be used inappropriately.

- Misunderstanding can arise either from badly constructed web pages, or from data that is not adequately processed for the level of decision making (e.g. information overload through the inclusion of too much raw or irrelevant data). These issues can be addressed by good design of web pages and good use of access control.
- Data from a remote web-enabled source may be used directly in a local control loop. This is not generally recommended by manufacturers [Babb (2000b)].

3.2.3 Data relocation

The relocation of data or expertise from the shop floor or from its point of application is facilitated by the use of internet and network technology. For many years the use of call centres has been increasing for information services (e.g. customer support lines, BT Red Care burglar alarm monitoring, fault reporting etc). This is the use of existing telecommunications networks for centralisation of skills and resources. These systems generally assume that the skills reside in the remote facility; that the caller can describe the symptoms appropriately but lacks the skill to diagnose the problem; that the caller has the skill to translate the facility's recommendations into the necessary local action.

There is scope with this centralisation framework to make use of internet technology. Where a person-to-person enquiry is needed, there is scope for use of voice and video as well as data over the Internet, and email may be practical. Where remote monitoring and control of plant is needed (as in the BT Red Care burglar alarm monitoring), query and response through a web page may be practical.

Data relocation separates the diagnosis expert from the site of the enquiry, and can remove some of the immediacy of the situation. The possibility arises that the expert's judgement may be influenced.

Separation from the site may prevent the expert from considering alternative diagnoses for the observed symptoms.

- In remote medical diagnosis [Jourdan (2000)] notes that the scheme had a consultant performing the diagnosis via ISDN video phone, but with the local GP in attendance with the patient. This provides extra diagnosis and observation skill at the point of application. If the GP were not present, the consultant would have to diagnose using vision alone and might miss some symptoms.
- In automated factory software, products such as Schneider Electric's "Transparent Factory" offer facilities to monitor and control a factory, and to provide remote maintenance and support. The remote system will only tell you symptoms of what it is capable of measuring i.e. those that have been foreseen as informative. The remote expert can view only a pre-determined set of diagnostics, where an expert on site would be in a position to notice something pertinent (e.g. a constricted pipe due to a misplaced crate). The remote expert's incomplete view can lead to unsafe decisions that may not properly address the cause of the problem.

The use of a video link for the remote engineer to gather contextual visual information, as well as perhaps a verbal description from the calling operator could help in reducing the possibility of a poor safety-related decision.

Where a semi-skilled operator on site reports symptoms to a remote expert, and the expert recommends safety actions, there is potential for miscommunication.

- The semi-skilled operator may not be able to observe and report the diagnostics in an appropriate form for communication to the expert.
- The remote expert may not be able to communicate the actions or advice at a level to suit the semi-skilled operative.

The communication between a local semi-skilled operator and a remote expert could take place over telephone, via WWW, or by email (if the response were not needed urgently). The content of the communication would require careful construction. A balance must be struck between (1) a text or free-form message that relies heavily on the observations and communication skills of the operator, and (2) a prepared set of "multiple choice" scenarios that may not cover accurately the particular scenario that has occurred.

Consider a system that uses mixed communications methods such as telephone for emergencies, WWW for urgent help, and email for general support. If this is intended for use by semi-skilled operators in a factory then clear definitions of what constitutes an "emergency", "urgent" or "support" would have to be given, else the system relies on interpretation of level, and the potential for confusion arises. If the internet/network facility were to be used, then clearly that system would have to be engineered to be fail-safe, and as part of the safety-related decision making process it would have to guarantee a level of service in a potentially harsh situation.

The addition of Voice over IP (VoIP), and integrated video, internet and phone systems is a potential failure to deliver emergency or safety-related data. As with modern small office telephone exchanges that are powered from the standard office power supply, there is no telephone service during a power cut. An integrated data and phone system puts all communication channels through one medium, and introduces the potential for a single-point failure.

The above scenarios consider the human workforce is still in place. As technology progresses towards full automation, the aim is for "lights out" factories with no human workforce [O'Brien (2000)]. Under these circumstances the emphasis changes from harm to the workforce, and instead considers the effects of mishap on the environment. The fully automated factory under

remote maintenance is an extreme example of the remote expert relying on the diagnostics at the plant. While remote diagnostics is addressed in HSE (1995), the element considered here is the reliance on network technologies (phone, ethernet (say), and associated internet protocols).

All the security issues discussed previously apply also to the web-enabled automated factory.

In the offshore industry there is much discussion [OILC (1999)] about "deskilling" and "multi-skilling" (where one operator covers many roles). The main concern here is the greater reliance on semi-skilled workers to perform tasks outside their capability. OILC (1999) gives some examples of Health and Safety breaches as an alleged consequence of deskilling.

All the issues discussed previously on the subject of remote expertise apply in the offshore scenario, where an internet-technology based application replaces a key service on a remote platform.

3.2.4 Technology uptake

There are various large systems currently in operation that use a mix of internet and more traditional control technologies. A wide area example is detailed in Pons (2000) which describes the remote control and operation of a water distribution system. This system has various considerations for the support of safety: the segregation of control and supervision; using different messages so to avoid ambiguity; the use of dedicated phone lines, and a star network topology so if a branch of the network fails then the rest can continue to operate.

A recent report by O'Brien (2000) describes a geographically smaller venture by Pirelli in the production of a totally automated factory. This factory can have production details altered from the central factory computer via a link to the internet.

When individuals are given access to data from all detailed aspects of (for example) a factory, a bewildering amount of data becomes available. It becomes increasing possible to make a poor safety-related decision with a wealth of irrelevant information. In addition to managing the volume of data, the presentation of the data is also important. For example the product "WizCon" from eMation uses representations of the installation with a strong graphical leaning (e.g. shadowing on pipes, pretend brickwork, shadows projected onto walls). In some circumstances this irrelevant graphical detail may lead to a misunderstanding of what is being observed, and then to poor input to a decision making process.

Among others, Rockwell and Schneider have incorporated web-servers into their PLCs [Prouty (2000)]. It is unclear from casual enquiries of vendors what the true market penetration of internet enabled PLCs or control systems is - this sort of information is in the sales regime and is considered confidential. Several industry sources claim that this web-enabling of controllers and diagnostics is being driven by customer demand, but hard evidence is scarce and it would seem prudent to assume a significant element of technology push by the suppliers to promote internet-based applications.

The Schneider web site for their product "Transparent Factory" does give access to a page detailing applications of the technology (www.transparentfactory.com). This applications page details 32 examples in use world-wide. They range in scope from candy making processes (which are directly linked to the corporate ethernet) to an offshore heavy lifting system which uses wireless networking and ethernet in the solution. With the latter example the system is not currently linked into the corporate network, though as has been commented earlier, networks may start off well-designed, but will get "added to" later in their lifecycle.

Further inquiries about Fieldbus and industrial network and ethernet connectivity showed that the area has not yet matured, and is very much in the prototype stage.

A recent development in safety systems uses a version of PROFIbus called PROFISafe. Little information was found about this system, though ICS (2000) gives general information, and some indication of the redundancy in the system. PROFISafe is intended to replace hardwired safety systems with a PROFIbus based one. This is not necessarily an unsafe way of working. However, future developments in connecting safety-related PROFIbus and Fieldbus to office networks or to the Internet have the potential to cause an unacceptable level of safety system compromise.

The connection of Fieldbus (PROFIbus) to ethernet through a gateway is still in an early phase of development. Tellima Technology offer a product ("Field Marshal") that performs this connection, but it is still in the prototyping phase. The connection of Fieldbus to ethernet was currently under consideration by a PROFIbus committee, as they were aware of the security implications.

It is not clear from discussions with suppliers just how much interaction is possible across the Fieldbus/ethernet gateway. But even if only limited interaction is possible, there will exist the opportunity to cause disruption to the connected Fieldbus either through mistake or malicious intent. As Fieldbus can be used in time-critical applications, missing a message by milliseconds may be crucial. There is a potential loss of messages from the ethernet local area network.

The two systems of (1) PROFISafe and (2) linking industrial network to ethernet have their own individual problems, and coupling the two could potentially double the safety-related issues. For example, if the PROFISafe system is performing a control action, and then a flood of commands enters the bus from the ethernet at the same time as a safety interlock trips, how will the combined system react? Can the PROFISafe system react in time if there are many commands in the system communicating with individual devices?

As detail of the planned rollout of Fieldbus on ethernet is readily available only to Fieldbus consortium members, this report will speculate no further on the state of development of this technology.

3.3 INTERNET TECHNOLOGY IN DOMESTIC LIFE

The influence of the Internet has so far only stayed at the "family browsing and email" level. Technology developments will add internet capability into everyday products. The domestic situations of concern have strong parallels with some industrial situations.

3.3.1 Domestic implications

It is currently possible to buy household goods that are web-enabled [Schofield (1999)] such as washing machines. It is not at all clear what are the practical benefits of this capability and why it might become attractive to the market. This is evidently an area where suppliers are pushing technology rather than consumers demanding it.

Future developments of internet technology will lead to convergence between the Internet, phone, mobile phone, and television networks. With technologies such as Bluetooth [Friedli (2000)] most household devices could be capable of wireless intercommunication. The benefits (according to the technology suppliers) are that it becomes possible to remotely control domestic operations such as turning ovens, fridges, washing machines etc on and off remotely. Clearly, remotely turning on an oven or a washing machine could have serious safety consequences as the local conditions cannot (potentially) be checked prior to engaging the device.

Another supplier, Pace Microsystems, has announced a partnership with Microsoft to support control via a "Home Gateway" which is networked through the television set-top box. This too

is described as being capable of controlling home appliances [Hayward (2000)], although significant further practical development is needed.

The main potential concern with this aspect of internet technology usage is in the remote operation of domestic devices, and the fact that there is potentially a massive installed consumer base. If web-enabled domestic equipment becomes common, then statistically there is a greater potential for device failure through malicious attack (see below), or accidental operation, than currently where device operation is a manual task.

Home users are as much at the mercy of misinformation as commercial users. Home users are more likely to seek medical advice, and may also seek guidance notes for products from the WWW. Due to the vast number of individual users, and the wide range of capabilities of those users, any information taken from websites will have to be exceptionally clear, to avoid a large number of poor decisions being made.

As telecommunications starts to integrate with Internet and television, a large number of services are going to use the same physical channel for transmission to the domestic user. This introduces a single point failure into the system, and may result in loss of emergency communication.

3.3.2 Industrial implications

Telecommunications technologies like ADSL are emerging that will give permanent unmetered internet connection to home users. This will increase the likelihood of malicious attacks into domestic systems, either conventional PCs or controllers of web-enabled domestic equipment. The domestic equipment controller is effectively a low capacity networked computer that is open to attack in the same way as the networked plant controllers discussed earlier in this report. Most domestic users are unlikely to have the awareness or the skill to install effective security measures to prevent unauthorised access.

This raises a possibility that may attract the attention of resourceful but irresponsible “hackers”, and one that is being actively discussed in internet newsgroups (e.g. comp.security, comp.security.misc, and comp.security.firewalls). It is technically possible to install on a suitable networked controller a program called a “trojan” which remains inactive until activated by predetermined conditions. It would be possible for a large number of trojans to generate a very large amount of network traffic that would maliciously disrupt a safety-related industrial activity that depends on timely net response. More generally, these home systems may make it easier for hackers to “bounce off” other peoples’ systems to cover their tracks. The web-enabled washing machine may be a highly implausible domestic novelty today, but its consequences could be serious.

3.4 DATA PROTECTION

Any aspect of business that relies on connection to the internet will require suitable protection whether it is in the distribution of pertinent, timely data, or the use of internet technology as part of a business or production process.

Network data security is often not a primary safety consideration, and has not traditionally been regarded as a safety issue for engineering attention. The traditional view is that if the engineers competently build a safety-related system, and if someone then maliciously sabotages that system, then the problem is one of law enforcement rather than engineering. However, a safety-related system should be designed with adequate attention to foreseeing and removing obvious vulnerabilities. A distributed network-based safety-related system is clearly vulnerable to data corruption, and effective data security mechanisms should be implemented to achieve adequate safety integrity.

3.4.1 General Security

There are many aspects to securing a network against attack either from a single protagonist, or from non-validated software intrusion. There are many products available to offer differing levels of security from all the many points of entry into a network. One aspect of security that can often be missed in the discussion of technology is simple social education.

If an organisation has adopted the Open Network model (see Fig.1) then the following sections should be taken into consideration. If plant controllers and similar safety-related devices are connected direct to the internet, then their security capability needs to be ascertained.

3.4.1.1 Social Education

This aspect is a keystone to any data protection policy. No matter what technology is in place for protection, active malpractice or casual neglect of standard procedure in the area of security can cause a bridgehead into a network.

Social education encompasses the way people use IT. General points to consider when trying to retain control of the network, and the integrity of its data include:

- Do not write passwords down as an aide-memoire.
- Do not throw away old passwords that have been written down (this can give guidance to intruders if someone uses a pattern to their passwords).
- Do not discuss network details outside the office.
- Do not allow SOUPs to run on browsers and email clients (this prevents Active Scripting, and ActiveX).
- Enforce strict use of security tools such as virus checkers.
- Do not download and run SOUPs such as screen-savers, games, utilities from the WWW.
- Scan all transport media (floppy disks, CDs) for viruses prior to use.

3.4.1.2 Virus Checking Policy

There are many products currently available that can protect against virus incursion, including the Active Scripting type of virus that caused the “Lovebug” incident. These can be installed on mail servers to screen incoming mail at the point of entry to the network, and can also be installed on individual desktop PCs in the network. Note that many ISPs also have a policy of scanning email before delivering it.

The main problem is that the effectiveness of a virus checker always lags behind new viruses, and the checker is ineffective against a new virus until it has been updated with the “signature” by which it can recognise the virus. An organisation should therefore not rely exclusively on virus checking, but should also implement the social education aspects in order to avoid the viruses in the first place.

3.4.1.3 Firewalls

This is a class of program that sits either physically or logically between a user and the internet. A firewall implements the organisation’s security policies: what type of electronic traffic is allowed; from/to which websites; the conditions on which material can be placed on or removed from the organisation’s computers.

3.4.2 Site Integrity

Web-site integrity covers topics such as security of the site, and the correctness and timeliness of data.

In order to be assured that the infrastructure on which the safety-related data resides is as secure as possible, an organisation needs to review which operating system and which web-server applications to use. This review could cover, for example, which systems have the best record for information availability. The review should also consider the choice of ISP with which to connect to the WWW: the ISP's technical characteristics and track record of performance.

This is clearly less of an issue for internal systems (i.e. "intranet"), though consideration should be given to having a diverse approach to office systems, controls systems and data distribution systems [Kerstetter (2000)] so that they are less likely to fall foul of the same problems at the same time.

The actual site content is also an integrity concern. The data should be correct. For example, all hyperlinks should work i.e. should point to a currently valid source of data. An auditing tool such as Microsoft Site Analyst should be used to audit the site content and to check the links to ensure the information pointed to exists.

A more technology-based check is the use of a port scanner (e.g. see (http://www.grc.com). All computers have connections to networks via "ports", and various functions of the computer use these ports to operate. Improved security results from limiting a function's access to those ports that it actually needs. For example, a function that accepts an incoming link from the Internet is unlikely to need "File" and "Print Sharing" enabled. A port scanner would list all functions on the computer that were offering such a way into the system, thus making possible a precise management policy on the privileges permitted to each function.

3.4.3 Network Topology

The choice of network topology and connection to the Internet will have an effect on the integrity of the data and infrastructure in use. Lafave (2000) discusses topology in detail in the context of single-network structures. Note the importance in such a network of single points of failure such as hubs or switches. Clearly if a safety-related system is to be hosted across a network then plans have to be in place to cope with hardware failure.

If connection to the internet is via a permanent link, then it is more likely to receive potentially malicious attention from outside [Cheswick (1994)]. To maximise security, when choosing the internet network topology, one should consider isolating internet based activities in a separate network section (i.e. "broken network").

If connection is via ISDN or via a dial-up modem to an ISP, then the connection does not exist permanently. In this case, the address of the computer on the Internet is dynamically allocated (DHCP - Dynamic Host Configuration Protocol) each time the connection is made. This makes the computer harder to find in the internet, and offers some protection against unauthorised access.

Whatever the connection between the internet and the LAN, there is a general consensus that keeping office network activities separate from control and automation networks is a good thing [Mintchell(2000)]

4. MANAGEMENT AND PROCESS CHANGE

4.1 DISTRIBUTED MANAGEMENT

4.1.1 What is this?

This term reflects the trend for large organisations to collaborate in the achievement of a goal, such as a project, or operational undertaking. For a non-internet related example consider the way British Rail has been broken up, so that the train network is now run by a collaborative effort of several companies, between them covering all aspects of train operation and infrastructure maintenance.

Distributed management structures may be adopted for a variety of reasons: privatisation break-up of organisations, preferences for project management, perceived cost advantages. Overall responsibility rests with multiple companies, and it could be perceived that this flattens the overall management structure. It also makes it difficult to apportion responsibility when things go wrong.

Within the public eye the Internet has a high profile, and companies seen to be embracing such technologies are perceived to be better than more traditionally managed companies. Also there is a lot of hype around the topic of internet technology use, and the perceived benefits (i.e. faster communications, secure transactions, business to business e-commerce, savings to be made, etc.).

On a lower level of organisation and control, there is the area of Highly Distributed Control (HDC) systems, which basically uses the embedded controllers as an intranet of web-servers [Prouty (2000)]. This arrangement makes the availability of data higher, and also to a higher level within the organisation.

4.1.2 Technology uptake

From the publicly available trade literature it is not clear that any companies have deliberately broken themselves up to take advantage of internet technology as part of their business infrastructure. What is happening though (from observation of the trade news and the Tessella customer base) is that companies are reducing their staff numbers as benefits are realised from more effective communications use of email, WWW and intranet technologies.

An example would be a large petrochemical organisation which has a scientific software department that is run from Holland, and has offices in various sites across Europe. With the effective use of email, phone and WWW the daily communication is uninterrupted, with only the occasional face-to-face meeting required, and the department runs efficiently. However, this may introduce the problems reported in [NASA (2000)], which include the difficulty (due to geographic or regional differences) of effectively communicating data that is described in a different system.

Another indicator of how much management is interested in what the plant, or manufacturing process is doing can be gleaned from the type of proposals recently undertaken by Tessella. The areas of Enterprise Application Integration (EAI), and importing plant control information into Enterprise Resource Planning (ERP), or Management Information Systems (MIS) is becoming more frequent. One system we inspected was designed to provide scientists with a wealth of unstructured data. This system was, in Tessella's opinion, a mistake as the scientist making the remote enquiry had to be proficient in a database query language such as SQL in order to extract data. This puts the enquirer too close to the data and returns an inordinate amount of unstructured data with no context. Any decision taken using data from this system is liable to be poor, because the enquirer would have too much irrelevant data and no context in which to

place it. What this indicates is the need for a query tool that produces pertinent data for each specific user-type.

One area that is mentioned in literature, and which Tessella has some experience of is in the area of linking "shop floor to top floor" together. This marketing phrase is used to described systems where the manufacturing plant data can be observed in the systems of the higher management levels. This allows, for example, a stock purchasing system to find out automatically from the plant or process whether new materials should be ordered. This area is also being heavily pushed by suppliers such as Rockwell, Schneider, and ERP system suppliers [Peach (2000)]. The use of internet technology to get data from "shop floor to top floor" has been quoted as helping organisations in meeting supply deadlines [Coates (2000)], but there is little data publicly available on specific uptake for safety-related systems, as this is usually confidential.

4.1.3 Examples of co-operative working.

"Shop floor to top floor" integration does not in itself lead to a health and safety issue, but does introduce the concept of "business to business" e-commerce, where an organisation makes use of distributed facilities to operate. For example, a business might arrange with a supplier to purchase goods via an electronic transaction. There is a perception that savings can be made in this approach, and this perception is forcing uptake.

If this model is extended further to purchasing services as well as goods, then safety issues may be raised. For example, if a chemical manufacturing plant is currently processing a particular chemical reagent, then it may order the appropriate specialist disposal service from a supplier to dispose of the waste products. As the order is electronic there may be no manual verification (at either end) that the correct service has been ordered. This could lead to the unsafe disposal of materials.

For large undertakings such as space missions there is a lot of co-operative working. Responsibility lies with the overall project manager, but the very nature of the global working introduces problems. The report NASA (2000) demonstrates how working across national and cultural boundaries can present a new set of difficulties that have to be understood. In the NASA (2000) case this was differing measurement units. If a safety-related decision has to be made based on data that crosses an abstract boundary (e.g. from Imperial to Metric measuring units, or differing acceptable doses in different countries) then the chance of a poor decision being made is high, unless the system is thoroughly designed and end-to-end checking is made possible.

Consider also the example of a local authority that has overall responsibility for emergency planning in a catchment area that contains several high-hazard installations. Close and effective co-operation is clearly needed, and the flexibility of an internet-based approach may well be a disadvantage, unless adequate management procedures are in place to address safety concerns.

4.2 DISTRIBUTED WORKING

4.2.1 What is this?

Just as overall management and control functions can be distributed across an organisation, there are possibilities for the working methods to change also. Distributed working may be the use of networks of freelance programmers [Lewis (2000)], or the movement of specialists from an office basis to a "roving" role. There has always been a home worker population, but with internet connectivity more workers may now be able to fulfil their work roles without leaving the home.

The use of internet technology to further distributed working will generally have a greater effect on the office worker. This class of worker uses IT as a communication and document preparation tool, so the use of an office's facilities may be superfluous. This raises the issue of the home as a work place. There may be health and safety implications, such as lone working, work conditions etc. Who is responsible for enforcing any regulations?

The worker that needs to tend a plant process will in general need to still need to be next to the machine they tend. Supervisors at factories, or maintenance engineers, specialist roles that do not require 100% attention on site can be moved over to a distributed work form, as one supervisor may be able to tend to two sites. This type of approach can be seen in companies that "out-source" key roles to third parties who then perform routines maintenance tasks on a rota basis.

4.2.2 Technology uptake

There are many products currently available which promote collaborative working, such as computer-aided design (CAD) packages that allow two designers to work on a project simultaneously. This approach to work productivity tools leads to the possibility of home working as an acceptable norm.

Other technologies that can allow a "roving" worker to access data include mobile and WAP (wireless application Protocol) phones. WAP phones allow access to the Internet, and so potentially any data that is available through the WWW can be accessed from a phone.

One potential vulnerability of these devices is being in an inappropriate location to access safety-related data, or to make a decision on that data. For example the operator of the phone may be about to submit a command sequence to a remote process when the train they are on goes into a tunnel causing loss of service. This indicates that a control process from an internet device needs a guaranteed quality of communications service. See [Lafave (2000)]

Another difficulty arises because using the smaller input devices on WAP phones is awkward, and there is a great potential for incorrect data entry.

An example of control based on mobile telephony is the Mitec (www.mitec.se) product that enables a remote sensor on a tank (such as an oil tank) to send a message about levels and flows to a web site or to a mobile phone through the Short Message Service (SMS). This type of technology allows maintenance workers to change roles from site-based to being able to monitor and assess a range of sites, so they become distributed workers.

4.2.3 Examples of distributed working.

One of the most clear changes in working policy over the recent years has been in the emergency services. Each health region now has a mobile doctor service, whereas doctors used to be more surgery-based. The use of networks by the London Ambulance Service has enabled a proportion of ambulances to be stationed not at ambulance stations [LAS (2000)], but instead at "standby points" where they are closer to the point of application than before.

One potential vulnerability of safety-related decisions made by a distributed workforce is that of monitoring the "teleworkers". How can it be verified that they are following correct procedure, and are suitably motivated to perform their duties in a timely fashion as they would if they were at a place of work? Newman (1997) identified discipline and motivation as two of the biggest problems of teleworking. The adherence to due process and providing information or decisions in a timely fashion is needed for any system that works with safety-related issues.

The IT service industry is also an area where a lot of teleworking can occur. A new industry is emerging in supplying "virtual project teams" for IT projects [Lewis (2000)]. Beyond the independent contracting field, large corporations such as BP have been using "virtual team

networks" for a while [Lewis (2000)]. The main issue here is to identify what is the work place, and who is responsible for health and safety.

With mobile network technology such as WAP phones there will be increased likelihood of people using these devices in inappropriate circumstances. For example, a current social issue is to do with people using mobile phones whilst driving. Will this worsen to people reading email whilst on the move? When email becomes available on a greater range of mobile devices, then the issue of distraction from a safety-related task may have to be considered - there will be increased opportunity to be distracted.

Another aspect of mobile computing is that a mobile terminal can access a computing network from anywhere, and may have been personalised by the user to provide rapid access to network data and resources. There was a celebrated case of an MI5 officer who had his laptop stolen from his car. If this had been a mobile network enabled device, then that organisation's security could have been breached, and access gained to network resources including potentially plant control.

In the scenario of a mobile engineer with remote access to a manufacturing plant from a networked laptop, the loss of this networked device could be a serious threat to health and safety. An internet-based control system may need some way of discriminating between trusted and untrusted remote terminals. It may be that mobile terminals are always untrusted as they are always potentially being misused, or used in an inappropriate place (on a train about to go into a tunnel...).

5. SUMMARY

There is a great deal of pressure for organisations to take up the internet tools and get themselves "web-enabled". There is outside pressure from the customers, and pressure from within an organisation.

Use of the World Wide Web influences how an organisation is perceived. It is becoming more frequent to ask for a URL in order to get further information. The WWW is very much in the public eye, so there is a strong move to use this customer awareness to publicise an organisation, and provide increased availability of information; this leads to an improved *perceived* level of service [Reeve (2000)].

Another source of pressure for the take up of internet based technology is from within an organisation. Ethernet and TCP/IP are well known and trusted networking technologies [Lock (2000)], and there is a perception that getting control systems onto Local Area Networks could save money [Reeve (2000)], and also provide more timely data for various departments such as Customer Relations. With the general push to cut overheads, the use of internet technology appears attractive [Clutton(2000)].

As well as the perceived benefits of connecting a business to the Internet, there are also disadvantages. Theses include: the well-known and well-publicised threats from viruses and hackers; potential inaccuracies in, and liability for, publishing data on which safety-related decisions are being made; exposing internal systems (business or control) to the outside world.

There are many manufacturers who are pushing the technology to allow connection of manufacturing and control systems to the Internet. There is no real evidence that adequate consideration has been given to the security or stability of these devices and systems, either to their ability to resist attack, or to their allowing unnecessary cross-talk into any time-critical manufacturing control networks. However, ethernet has been used for many years in an office environment and is a well-known and trusted technology. From the consumer's point of view, this de-mystifies network-based automation and control, and makes the approach more acceptable, although not necessarily any better understood in reality.

Introducing a "shop floor to top floor" integrated system to allow manufacturing data into the business management systems would have to be done in a very controlled way. There is scope for information overload (of management), and for making decisions on irrelevant, or raw, data.

The safety considerations of allowing a direct connection of a controller to an ethernet local area network, or of connecting a control network to an ethernet local area network via a gateway, are still under consideration by such organisations as PROFIbus. Hardware does appear to be available to perform this connection, but the resilience of the configured system to attack, or being swamped with irrelevant data is unclear. This area should be looked into very carefully for potential safety issues. There may need to be some form of certification for any device that allows access to control or safety systems. This certification may be a declaration that the system has been tested for operation under representative and demanding circumstances, such as those detailed by CERT in their advisory reports on system weaknesses. This will be of significant interest with the advent of PROFIsafe, the integration of safety control with process control on the same controlling system.

The Graphical User Interfaces used to display the data for engineers performing remote operations varies greatly from product to product. A sophisticated graphical display may contain much irrelevant and potentially confusing detail, and may detract from the safety of a decision. There might be benefit in considering how Human-Machine Interfaces such as GUIs

may be standardised by an industry body in the same way that technology interfaces are agreed by committees.

It has been noted that the quality of published safety-related data varies. Examples have been found that are potentially dangerous, while other examples demonstrate good practice by including all the relevant references and guidelines. Where a company publishes guidance for health-related products, good practice is to review this information regularly. A related question to be resolved is that of the responsibility for distributing safety-related information. Issues highlighted include the caching of data on web servers and web browsers, and the upkeep and maintenance of the website and links to and from that site.

Network integrity is paramount in any internet-based system that is to distribute safety-related information, or for control and automation. Some sites of high profile are heavily targeted by hackers. A high profile results from (1) traditional publicity surrounding the organisation in question, and (2) the visibility of the site on the Internet. The Pentagon is a particular target, but with growing emphasis on ecology issues in public life, any companies associated with bad publicity to do with the environment are also likely to become targeted. Consider the scenario of a chemical company that has recently been publicised as harming the environment. In response there may be denial of service attacks from hackers, or even attempts at intrusions. If the network under attack is also connected to control systems, this could have serious safety repercussions. Various methods of protection were discussed, but the best advice is to separate the networks for the plant control and business infrastructure [Mintchell (2000)].

Material on where the use of the Internet or internet technology has changed management or working styles is difficult to locate. What is easier to locate is management structures where implementing an Internet related solution would be inappropriate. Lack of access to pertinent safety data, no definite co-ordinator (hence difficult to attribute responsibility), a need to use data in a mobile environment (such as needing to read procedures in a plant in a hazardous environment) are issues that need to be considered before implementing a management system based on internet technology.

The use of internet technology based systems can easily give rise to unsafe “systems of work”. This is mainly because the procedures are designed to optimise efficiency, but can contain serious failure possibilities. Available internet technology is such that a safety-related process can be moved to an internet-based solution in a very short time. Adequate planning and management is needed to ensure that no essential safety aspects of the new process are overlooked.

REFERENCES

Babb, M. (Ed.), Wright, J. (Pub.) (2000a)	'A yellow card for IEC 61158', Control Engineering Europe, April 2000, pp24-25.
Babb, M. (Ed.), Wright, J. (Pub.) (2000b)	'PLCs on the Web', Control Engineering Europe, April 2000, pp29-30.
Braden, R., Zhang, L., Besson, S., Herzog, S., Jamin, S. (1997)	RFC 2205 [online], Connected: An Internet Encyclopedia, Available from: http:// www.freesoft.org/CIE/RFC/2205/index.htm [Accessed 01/06/2000]
CERT	Computer Emergency Response Team, website (www.cert.org)
Cheswick, W.R., Bellovin S.M. (1994)	Firewalls and Internet Security: Repelling the Wily Hacker, Addison-Wesley
Chlor-Chemicals (2000)	Typical properties of Trichloroethylene [online], Chlor-Chemicals Ltd., Available from: http://www.chlor-chemicals.com [Accessed 18/5/2000]
Clutton, A. (Ed.)	'Working benefits of the Web', Plant & Control Engineering, February 2000, pp26-27
Coates, T. (2000)	'Integrate below the ERP level', Control and Instrumentation, January 2000, pp21-22
EPA (2000)	OAR Policy and Guidance Metarecord [online], Environmental Protection Agency (US), Available from: http://www.epa.gov/ttncaaa1/t5/meta/ m16796.html [Accessed: 18/5/2000]
Everlac (2000)	Epoxy Floor Paint [online], Everlac, Available from: http://www.everlac.co.uk/health.htm [Accessed: 18/05/2000]
Friedli, D. (2000)	'Domestic appliances get connected', The Engineer, 21 April 2000, pp18-20
George, W. (Ed.) (2000)	'Where Process Control meets the Internet', Industrial Networking + Open Control, Volume 6, April 2000, p11.
Gray, D.F. (2000)	German court: AOL responsible for pirated music [online] CNN, Available from: http://www.cnn.com/2000/TECH/computing/04/14 /aol.pirates.idg/index.html [Accessed 01/06/2000]
Hayward, D. (Ed.) (2000)	'Business Web: Microsoft and Pace automate the home via the Net ', Computing, 11 May 2000, p4.
HSE (1995)	'Safety in the remote diagnosis of manufacturing plant and equipment', Health and Safety Executive, HS(G)87

HSE (2000)	Information Services [online], Health and Safety Executive, Available from: http://www.hse.gov.uk/org/orguk99.htm [Accessed 16/05/2000]
ICS (2000)	PROFIbus offers new safety solution [online] Instrumentation & Control Systems (PennNet), Available from: http://ics.pennnet.com/home/archivearticles.cfm?Section=Archive&ARTICLE_ID=65302&VERSION_NUM=1&KEYWORD=profisafe [Accessed 06/06/2000]
Jourdan, T. (2000)	Diagnosis by remote control [online], The Electronic Telegraph, Available from: http://www.telegraph.co.uk/et?ac=002575848746865&rtmo= InFzozut&atmo=ttttttd&pg=/et/00/1/11/thela11.html [Accessed 24/05/2000]
Kerstetter, J. (2000)	'Don't leave the field open to viruses', Computing, 25 May 2000, p26
Lafave, L. (2000)	'Digital Communication in Safety-related Systems' , Safety Systems Research Centre, University of Bristol, SSRC/HSE-1-1999, Draft Version 0.2
Laybond (2000)	Laybond CV [online], Laybond, Available from: http://www/laybond.co.uk/datasheets/cv.htm [Accessed 18/5/2000]
Lewis, J. (2000)	'Working together - separately', Computing, 1 June 2000, p49-51
Lock, G. (2000)	'Trapped in the Ethernet?', Industrial Networking + Open Control, Volume 6, April 2000,pp8-9
Mintchell, G. (2000)	'Open Control', Control Engineering Europe, April 2000, pp38-39.
Moody, G. (2000)	'The great Web threat', Computer Weekly, 24 February 2000, p47
NASA (2000)	Mars Climate Orbiter Failure Board Releases Report [online], National Aerospace and Space Administration, Available from: http://www.marspolarlander.net/msp98/news/mco991110.html, [Accessed: 22/05/2000]
Newman, D.R. (1997)	IT and work [online], Queens University Belfast, Available from: http://wwwparent.qub.ac.uk/mgt/itsoc/itwork.html [Accessed 25/05/2000]
O'Brien, L. (2000)	'Pirelli's grip on robot production', The Engineer, 28 April 2000, pp14-16

OILC (1999)	Don't shout at me! I'm the chef - he's the medic [online], OILC - The Offshore Union, Available from: http://www.oilc.org/BO58/BO58.deskilling.html [Accessed 25/05/2000]
Peach, M. (2000)	'Enterprise and controls companies converge on information management', Control Engineering Europe, April 2000, pp32-36.
Pons, M. (2000)	'Super Rimiez: Real-time control of a potable water network', Industrial Networking + Open Control, Volume 6, April 2000, P14.
Prouty, K (2000)	'Beyond the hype', The Engineer, 24 March 2000, p30
Randell, B., Benjamin, R., Gladman, B. (1999)	'Business at war', The Computer Bulletin, November 1999, pp23-27
Reeve, P (2000)	'Dive in and get WET!', Industrial Networking + Open Control, Volume 6, April 2000, pp4-5.
Reuters (2000)	French ISP seeks anti-Nazi help to monitor websites [online] Findlaw Legal, Available from: http://legalnews.findlaw.com/legalnews/s/20000525 /franceinternetnazi.html [Accessed 01/06/2000]
Schofield, J. (1999)	'Give your washing machine a Web page', Computer Weekly, 9 December 1999
Schulzrinne H., Casner, S., Fredereick, R., Jacobson, V. (1996)	RFC 1889 [online], Connected: An Internet Encyclopedia, Available from: http://www.freesoft.org/CIE/RFC/1889/index.htm [Accessed 01/06/2000]
Siemens (2000)	Process automation that leverages the internet [online], Siemens, Available from: http://www.procidia.com/whitepapers/leverage.asp [Accessed 15/05/2000]
The Chemical Company Limited (2000)	GRAF BOAT epoxy [online], The Chemical Company, Available from: http://www.holdich.demon.co.uk/chemical/ boat.htm [Accessed: 18/05/2000]
WDGH NHS Trust	Haematology Index [online], Dorset County Council, Available from: http://www.dch.org.uk/pathology/Haematol/Haematology_Index.htm [Accessed: 18/05/2000]
Wessex Resins (2000)	PRO-SET [online], Wessex Resins, Available from: http://www.wessex-resins.com/PRO-SET%20Pages/Pro-setFrameset.html [Accessed 18/5/2000]

GLOSSARY

ISP	Internet Service Provider
ADSL	Asynchronous Digital Subscriber Line
CERT	Computer emergency response team
COM	Component Object Model
DDOS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DOS	Denial of Service
EAI	Enterprise application integration
ERP	Enterprise resource planning
Firewall	Hardware or software protection against entry to a network
FTP	File Transfer Protocol
gopher	primarily a text document indexing and searching service
Intranet	Internal network implementation of web-servers
ISDN	Integrated Services Digital Network
ISP	Internet service provider
LAN	Local Area Network
MIS	Management information system
PLC	Programmable Logic Controller
PN	Private Network
Portal	Computer providing access to the Internet, or Internet services
QMS	Quality Management System
Router	Routes network data packets according to destination
RSVP	Resources Reservation Protocol
RTP	Real Time Protocol
SCADA	Supervisory Control and Data Acquisition
SOUP	Software Of Unknown Pedigree (unvalidated software)
SQL	Structured query language
URL	Universal Resource Locator (web address)
VBScript	Visual Basic program, embedded in a web page
Virus Scanner	Software to check all data for inclusions of viruses.
VoIP	Voice over IP
VPN	Virtual Private Network
WWW	World Wide Web

Printed and published by the Health and Safety Executive
C1.25 01/02

MAIL ORDER
HSE priced and free
publications are
available from:
HSE Books
PO Box 1999
Sudbury
Suffolk CO10 2WA
Tel: 01787 881165
Fax: 01787 313995
Website: www.hsebooks.co.uk

RETAIL
HSE priced publications
are available from booksellers

HEALTH AND SAFETY INFORMATION
HSE InfoLine
Tel: 08701 545500
Fax: 02920 859260
e-mail: hseinformationservices@natbrit.com
or write to:
HSE Information Services
Caerphilly Business Park
Caerphilly CF83 3GG

HSE website: www.hse.gov.uk

CRR 408

£10.00

ISBN 0-7176-2268-1

9 780717 622689