

A broadcaster's view of Security, interfaces and layers

Andy Leigh

EBU Project Group N/Security

This article focuses on the risks associated with introducing new communication and computing technologies into the broadcasting world. It is mostly concerned with inter-broadcaster connections, rather than issues that are internal to a specific organization.

1. Introduction

For decades broadcasters have been inter-working with each other, mostly using audio and video circuits. In some cases the circuits were one-to-one private arrangements and in other cases were supplied by common carriers such as the EBU. Generally, command-and-control signals – such as circuit set-up and tear-down – have not travelled on the same link as the media.

With the introduction of cheap and readily-available computer platforms and self-routing networks such as IP (which also carries its command-and-control signals in-band), broadcasters are expecting to reap economic, flexibility and efficiency benefits.

There is a price to pay for these gains. In most cases, the benefits outweigh the risks. However, there is a strong likelihood that organizations do not understand the risks they are taking. This article – based on a paper prepared for the EBU N/Security group – is an introduction to the issues and possible solutions.

1.1. *What makes a broadcaster different from other organizations?*

It is occasionally claimed that broadcasters and programme-makers are so radically different that the rules and practices developed in other industries and organizations cannot be made to apply in our field of work. In practice – as broadcasters become progressively more dependent on IT for making and distributing programme material, schedules, control-signals and metadata – the differences are mostly vanishing.

Since the similarities are so many, it is best to concentrate on the differences:

- **“Size” of data**

A piece of audio or video material is considerably larger than the kind of content that most businesses move across their infrastructure. However, it should be noted that control-signals, metadata, EPGs etc. are similar in size to most business data.

- **Clock criticality**
Many commercial organizations operate a “just-in-time” methodology, but few non-broadcasters need to hit timed junctions and – for real-time contribution and distribution – frame-perfect accuracy of timing.
- **A highly inter-connected “boundary-less” model**
Put simply, broadcasters have been inter-connecting using audio and video (analogue and digital) circuits for years. As will be shown later, this has never been a significant risk. In comparison, most other organizations do not have anything like the same level of open inter-connectedness.
- **The prioritisation of Confidentiality, Integrity and Availability**
Because it was initially concerned with securing military secrets on computers, many people assume that Information Security is wholly concerned with “keeping secrets” – known as **Confidentiality** (i.e. who can read or view the content or information). Broadcasters, unlike military installations and the bio-chemical industry, are probably *more* concerned about **Integrity** (who can create, change or destroy the information or content) and **Availability** (who can stop a service from being available to legitimate users). Trustworthiness is a fundamental concern for many broadcasters and it is paramount that they can be confident that what they transmit has not been tampered with. Similarly, whilst unexpected downtime is a bad thing for many companies, their failures are not instantly obvious to millions of customers in the same way that a transmission outage is. Consequently (with a few exceptions), broadcasters would normally rate Integrity and Availability ahead of Confidentiality.
- **Creative use of technology**
Broadcast and production staff are very creative and consequently are often more technologically aware than their counterparts in other industries. This dynamism can lead to them using technology in new and original ways which often benefit their company. It can also lead to increased risks when products are used outside of their design tolerances.

In all other regards, a broadcaster can be considered similar to most other organizations, in which case the risks that the broadcast industry faces and the solutions it needs to deploy are comparable to manufacturers, suppliers, shops, logistics companies etc. There is no good reason why a virus that affects a car maker should not disable a broadcaster. A hacker is likely to earn as much “respect” from gaining access to a broadcaster as he/she does from breaking into a well-known high-street chain-store. In fact, due to their high profile, and their trustworthiness, most broadcasters are likely to be bigger targets than many other areas of industry.

2. Background: problems and principles

In the 1970s experts began to publish papers highlighting their concerns at the quantity of sensitive information that was being stored and processed on shared computers and networks. Such computers, they reasoned, were not as strong as the locks and safes that had traditionally been used to stop sensitive information from being tampered with. The experts published a number of papers on the problem and suggested solutions. The principle that all these papers were considering was this:

“ Computer and network security issues stem from the act of sharing a processor (or memory, disk, store, wire, network, frequency etc.) with someone else. ”

One such paper was: “The Protection of Information in Computer Systems” by Jerome Saltzer & Michael Schroeder [1]. It stands out because it proposed eight principles which are still considered relevant 30 years later:

- 1) **Economy of mechanisms**
Keep the design as simple and small as possible. Complexity is the enemy of security and reliability. As a system becomes more complex, it becomes impossible to exhaustively test all of the functions and all of the potential failure mechanisms.

- 2) **Fail-safe defaults**
The default situation should be “all access is denied unless specifically approved”. Designs must be based on arguments around why an object should be accessible, rather than why it should not.
- 3) **Complete mediation**
Every access to every object should be checked for authority. This is one of the most fundamental principles to pursue when it comes to protecting systems, content and organizations. The implication is that every individual or system that requests access can be identified uniquely – which in many cases also implies that they are “logged in”.
- 4) **Open design**
The design should not be secret. This is often paraphrased as: “security through obscurity is not a security mechanism”. An analogy is the front-door lock; the way the mechanism works is not secret, but the actual shape of the key is.
- 5) **Separation of privilege**
Where feasible, two separate keys are better than one. This principle, which avoids security mistakes, is used in nuclear launch systems and also in banking systems (such as ATM maintenance).
- 6) **Least privilege**
Every program and every user of the system should operate using the least set of privileges necessary to complete the job. In other words – if a person needs something to get their job done, they should be given access to it, but everything else is inaccessible. This “principle of least privilege” has been shown to be one of the most valuable because it can prevent many incidents and accidents. It also leads to the standard advice: never do anything when logged in as “root” or “administrator” unless you absolutely have to.
- 7) **Least common mechanism**
Minimize the amount of mechanism common to more than one user and depended on by all users. Every shared facility represents a potential information path between users and must be designed with great care. Again, this is a critical principle to apply to systems design.
- 8) **Psychological acceptability**
It is essential that the human interface be designed for ease of use, so that users routinely and automatically apply the protection mechanisms correctly. Really good protection mechanisms are those that don’t interfere with the people using them. People become frustrated with bad mechanisms and they investigate workarounds. There is quite a bit of friction between this principle and the other seven.

3. Security risks and their influence on broadcasters

3.1. “E-criminals”

As the use of the internet has grown, it has become apparent that it is being used for criminal activities. No new fundamental criminal behaviour has been discovered – only online versions of crimes that have been around for millennia. There appears to be a fundamental link between “utility” (how many people can afford to have access) and the likelihood of criminal behaviour. Most criminals want to minimize their costs and risks and maximize their financial rewards and benefits, and in some cases the publicity they get. This has a significant impact on whether a system is attacked or not.

For example, if two organizations share a proprietary network, where the interfaces cost €500,000 each, the network is unlikely to be attacked. The expensive hardware and proprietary knowledge required make entry costs too high. With only two organizations on the link, the benefits are also likely to be too low.

Abbreviations

ATM	Asynchronous Transfer Mode	IT	Information Technology
ATM	Automatic Teller Machine	LAN	Local Area Network
CERT	(US) Computer Emergency Response Team	N/A	Not Applicable / Available
DMZ	De-Militarised Zone	OS	Operating System
DNS	Domain Name System	OSI	Open Systems Interface
EPG	Electronic Programme Guide	PDH	Plesiochronous Digital Hierarchy
FTP	File Transfer Protocol	PKI	Public Key Infrastructure
HTTP	HyperText Transfer Protocol	POP	Point Of Presence
ICMP	Internet Control Message Protocol	SNMP	Simple Network Management Protocol
ID	IDentification / IDentity / IDentifier	SSL	Secure Sockets Layer
IDS	Intrusion Detection System	TCP	Transmission Control Protocol
IP	Internet Protocol	UDP	User Datagram Protocol
IPSEC	IP SECurity	VPN	Virtual Private Network

On the other hand, if a million organizations share an open-standards network where the interfaces cost €200 each, the number of criminals is likely to be high. The knowledge and hardware entry costs are low and the large number of potential targets means that the benefit is likely to be great.

3.2. *The changing face of broadcasting and its impact on security*

For decades, broadcasters have sent signals to each other over expensive point-to-point or point-to-multipoint networks where the cost of entry has been high and the number of broadcasters on each segment of the network has been low. This is one reason why an attack against broadcast infrastructure would only have resulted in limited reward for the perpetrator. There is a possibility that this has led to endemic risk-complacency within the broadcast community when it comes to assessing the likelihood of criminal attacks against their infrastructure.

The introduction of highly popular, low-cost architectures such as PC/MAC platforms, Windows/Linux Operating Systems and IP networks for the movement and processing of broadcast content, heralds a future of improved flexibility, dynamism and value for money. The low equipment cost and popularity of the technologies also mean that broadcast infrastructures, and the content they carry, are now joining a world where attacks are economically viable.

3.3. *Computer network risks and solutions*

Traditional broadcast networks are based on telephony models where circuits are set-up, used and then torn-down. In comparison, at the lowest connection level, IP networks are always physically connected. Information transfer only happens when a source device places packets on the network and a destination device reads them off of the network. The IP packets only store the address of the destination machine and the address of the source machine in their headers. The packets record no information about the links that they need to traverse or the best way to get to the destination. Such decisions are taken by the network nodes which read the destination address and forward the packet down the link that they calculate to be the shortest or most appropriate.

This nature of IP routing means that every IP-connected machine can be totally accessible from any other IP connected machine. It also means that it is not possible for the network nodes to determine whether a connection is valid or not – the responsibility to validate the link lies with the receiver. The IP packet headers are not strongly protected and as a result they can be, and are, modified to hide the real source address. For these reasons, the information in an IP packet can never be trusted.

Therefore, to operate securely on an IP network, every device should be configured with fail-safe defaults [Principle 2 above] and reject all connections unless the packet can be shown to have the correct privileges and can be authenticated [Principles 6 & 3 above]. In practice, most computers

are not set up this way. The same is true for most “black-box” devices with built in IP technology. The only IP-level technology that can enable IP authentication and access-control is the “IPSEC” VPN (Virtual Private Network) solution [2].

Filtering firewalls are devices that inspect every packet and then forward the packets on, but only if the combinations of source and destination addresses match a set of pre-defined rules. Advanced filtering firewalls keep track of connections – ensuring that newly-arrived packets make sense in the context of what has gone before. Advanced filtering firewalls can also inspect and make decisions on the TCP and UDP connection addresses ¹.

Filtering firewalls can protect against many network-type attacks. But they have no strength against application attacks and hop-through attacks (see Sections 3.4. and 3.5.). A different sort of firewall called a “proxy” can be used in this case. Proxy firewalls are dual-network-card devices that listen on one network to incoming packets destined for a specific application (such as a web server). Rather than directly forwarding the incoming packets, the proxy firewall uses a special proxy application to generate fresh packets – based on the incoming request – onwards from the other network. As a result, only valid and well-formed application-level queries traverse the proxy firewall. The downside is that proxy firewalls tend to have lower throughput and higher latency (time delays) than filtering firewalls.

3.4. Application and server attacks

If we require an individual to access a certain application on a machine, then a filtering firewall will not be able to determine valid versus invalid access attempts. Although a proxy firewall might help, the only system that can detect if the incoming connection is behaving properly is the application itself. Robust applications perform a number of tests to ensure that all access requests are authenticated, valid and correctly formed [Principles 3 and 6 above]. Unfortunately, building robust applications takes time, money and skill. As a result, most commercial applications and many Operating System components are not robust. Consequently, a large number of recent attacks are targeted at weak applications.

Some computers don't just host applications. Some specialise as servers: computers that hand out web-pages, files, print-jobs etc. Technologies, such as web and FTP servers, which were devised for the Internet, are often relatively robust. File servers, which operate by letting the user “mount” the server's drive on their computer as if it were their own (the ubiquitous F: drive), were designed for internal LAN usage. As a result, mounted file serving – which is the way that most desktop production systems work – is extremely insecure and should not normally be allowed to connect to any non-trustable networks.

3.5. Hop-through attacks

Fig. 1 shows a typical security mistake.

Because the computer is not set up to forward packets from one network card to the other (i.e. it is not acting as a router) it is assumed that Principle 7 (least common mechanism) is being met because there is no common network between the unsafe and the safe networks.

Unfortunately, this is not true of the Operating System. Nor is it true of any network application – both are common mechanisms between the two network cards. An installation of this kind is only secure if the Operating System and all the applications on the computer have been specifically designed to be very robust and are kept patched.

1. TCP (Transmission Control Protocol) and UDP (User Datagram Protocol) are the multiplexing technologies that ride over the top of IP connections. TCP is connection-oriented and supports facilities like HTTP (TCP port 80) and POP3 (TCP port 110). UDP is connectionless and supports DNS (UDP port 53) and SNMP (UDP port 161). See [3] & [4].

To give an example of hopping from one computer to another, we must imagine that the computer above is based on a UNIX operating system and has both a telnet server and telnet client running on it². If an attacker telnets onto this box, they will be able to enter commands on this computer. If they enter the command “telnet” they will then be able to telnet from this computer to any other computer on the unsafe network and also the safe network. This works because telnet is a **common mechanism** between the safe and the unsafe network. In doing so, the attacker has hopped (at the application-level of the OSI stack – see section 7.) from **their** computer to the computer shown in Fig. 1. From there, they can hop onto any computer in the safe network.




Figure 1
A popular solution for joining two networks (e.g. a broadcast and business LAN)

These hop-through attacks are utilised by Internet Worms with devastating impact – rates of infection can exceed thousands of devices in a few minutes. Hop-through attacks can also be used by hackers to gain access to machines running on private networks. The technique is also used to hide the origin of the attacker.

Firewall managers are frequently asked to open a connection from the outside to a specific computer on the inside. The requester assumes that this specific connection will only enable access to the nominated machine. In practice, such a firewall-hole can easily enable a hop-through attack. The firewall manager should only agree to such a hole if he/she can be convinced that the nominated machine has no weak Operating System components or badly-written software on it, or that the computer itself is ring-fenced off (possibly using more firewalls) from all other devices.

3.6. Types of assailant

Analysis has shown that attacks can come from any number of quarters, depending on the context of the attack and the organization being attacked. The following list of adversaries has been collated by Bruce Schneier [5]:

- malicious insiders;
- organized crime;
- industrial espionage agents;
- hackers and virus/Trojan writers;
- lone criminals;
- terrorists;
- national intelligence agencies;
- infowarriors;
- the press;
- the police.

2. Telnet is an application that allows a remote individual somewhere on the network to open a command-line interface on the computer.

Obviously, not all of these would apply at any one time to a broadcaster, and some are more likely than others. It is however possible to imagine a situation where a broadcaster has run a story that has a negative effect on one of the groups listed above and where that broadcaster is consequently affected by a denial-of-service attack launched against them.

In practice, the highest volume of disruption is likely to come from viruses, worms and Trojans and, because these types of attack are not targeted at a specific person, organization or computer, the protection against them needs to be general and universal.

3.7. *Protecting the information AND protecting the systems*

In a very large network, there are far more people who have access to the edge systems than have access to the intermediate central nodes. The people that can access these central nodes (i.e. Internet Service Providers and telecommunications engineers) are normally well-trained and contractually bound to not attack the networks or the information that flows over them. As a result, attacks of data “in-flight” are fairly rare and relatively hard to do. Most attacks are initiated *from* an individual using (or software running on) a network-connected computer and are directed *at* an individual, organization or computer connected to the network. Occasionally the attacks are targeted at the central communication systems.

In an ideal world, all software would be perfectly written; all databases would be properly built; all Operating Systems would be flawless; all computer hardware would be professionally installed, managed to good security standards and kept up to date; all users and systems would be uniquely identified and authenticated using long and impossible-to-guess keys; all network transactions would be encrypted and cryptographically signed, and all systems would log all transactions and detect anomalies and self-heal.

As we have already shown, in the real world, most systems are flawed, badly installed and managed. Hack attacks and viruses/worms rarely depend on a guessed password. Instead they target a flaw in the Operating System that allows an unauthenticated person or worm to gain low-level access to a part of the computer. Inside the computer, the person or worm tinkers with the set-up of a badly written piece of software, which fails and allows the unauthenticated person or worm to become the system’s administrator.

So if two organizations exchanging data decide to encrypt it, they will manage to prevent attacks against the data in-flight. However as we see above, such attacks are relatively rare. These two organizations will also need to ensure that all the computer systems and network technology that they deploy is well built, up-to-date and well managed: otherwise, the encrypted content will not give them a great deal of security.

3.8. *The attack landscape*

Fig. 2 shows the increasing number of incidents reported to the US Computer Emergency Response Team up to 2003 [6].

Most of these incidents were attacks against computers and systems, not against data in-flight. Each incident involved anything from one site to hundreds (sometimes thousands) of sites. Some of the incidents had ongoing activity for long periods of time.

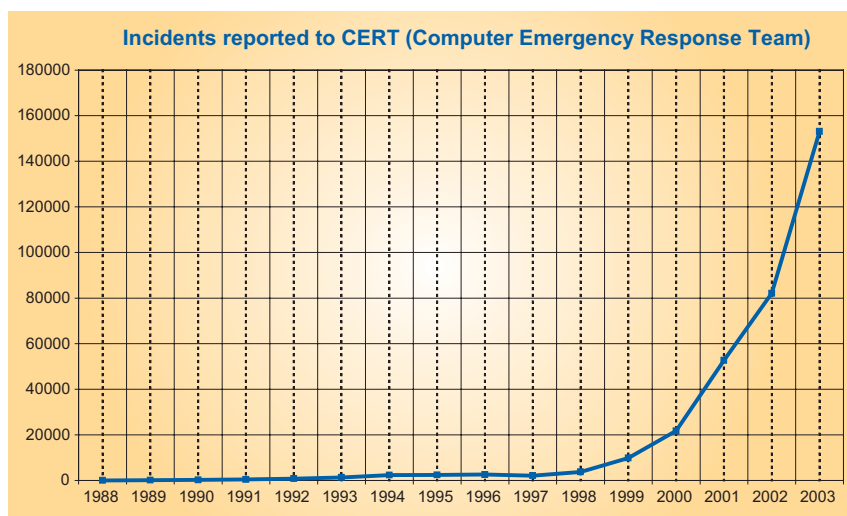


Figure 2
Increasing numbers of incidents reported to the CERT

4. The boundaries between organizations

One company does not normally connect itself to another without first agreeing the rules of engagement. In some cases there may be financial issues (e.g. payments for a service etc.); in other cases, there may be contractual issues (e.g. service levels and credits etc). In this section we consider the security interfaces between connected organizations.

Between any two organizations there is an interface. This interface might be physical, such as a door, or it might be logical, such as a contract. In the communications arena (which includes broadcast and computer networks), there is often an electrical interface or black box that enforces the policies agreed between the two parties. If we apply each of the eight principles described in *Section 2.*, we can infer that this black box should:

- be simple and easy to understand;
- by default, block all access to it and through it unless that access is expressly permitted;
- ensure that access to it, and transactions through it, are authenticated (i.e. performed by people or systems that have identified themselves uniquely and authenticated themselves, e.g. through a password certificate or key);
- be based on an open design – not secured through obscurity;
- have a completely locked-down configuration that can only be altered by both organizations with mutual consent (not by one of them alone) ³;
- restrict access to only those agreed facilities that are absolutely necessary but to no others;
- only utilise the bare minimum of shared apparatus between the two organizations;
- be easy to use or transparent to the people and systems that (i) need to use it, (ii) are authenticated and (iii) can gain access to the objects they need.

As stated in point (e) above, it is extremely difficult to build a single box and still maintain “separation of privilege”. In reality, every organization is responsible to its owners or shareholders for securing its own information. So if company “A” has a relationship with company “B”, company A is not empowered to fully trust company B, even if company B is large, wealthy and maintains good engineering practices. Each company builds its own “secure shell” or sets of interfaces that it controls to ensure that a hop-through attack (see *Section 3.5.*) does not affect it. The space between the shells is sometimes known as a **DMZ** (if it offers limited functionality) and at other times it is known as an **Extranet** (see *Fig. 3*).




Figure 3
The security interfaces between various organizations

- For one-to-one situations, this might be feasible but, for multiple interconnections, it is not. In practice, it is usually achieved by having two black boxes with a wire between them. One organization manages one box and the other organization manages the other box – the connection between the boxes is known as the DMZ (De-Militarised Zone).

To understand why companies do this, consider *Fig. 4* which shows what happens when Organization **C** suffers a security breach initiated on the Internet

In the above scenario, Organizations **A** and **D** as well as an intermediary (such as the EBU) will become totally dependent on their own interfaces to ensure that the attack does not hop-through and affect them.

Now consider what would have happened if (to ensure high throughput and low latency), there had been no security interface between the intermediary and its members – **B**, **C**, **D**. All of these organizations would have fallen victim to the attack as it hopped through from one infrastructure to the next.

This will have an effect on how the EBU considers setting security standards. Even if the EBU were to design and operate the best technology and best security practices available, many other companies could not and would not lower their defences in a total-trust relationship with the EBU. Even if the EBU can agree standards that are adopted by all broadcasters, any broadcaster will still wish to control its side of the relationship to ensure that it maintains “separation of privilege” (Principle 5).

There are possible operational standards that might emerge regarding “federated identity management” (which will enable Principle 3 – complete mediation) and transaction-based signing and encryption, but it will be many years before all organizations are able to participate enough for total and automatic trust to be invoked between any two organizations.

5. Relationships that could be considered

In the expanding world of interconnected organizations, the following types of relationship could be taken as part of the problem-space:

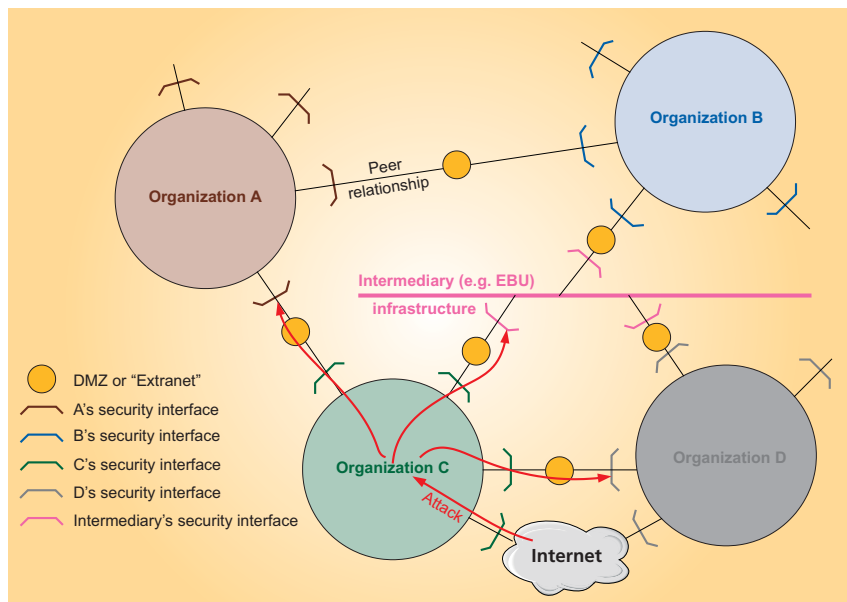


Figure 4
Organization “C” suffers an attack from the Internet

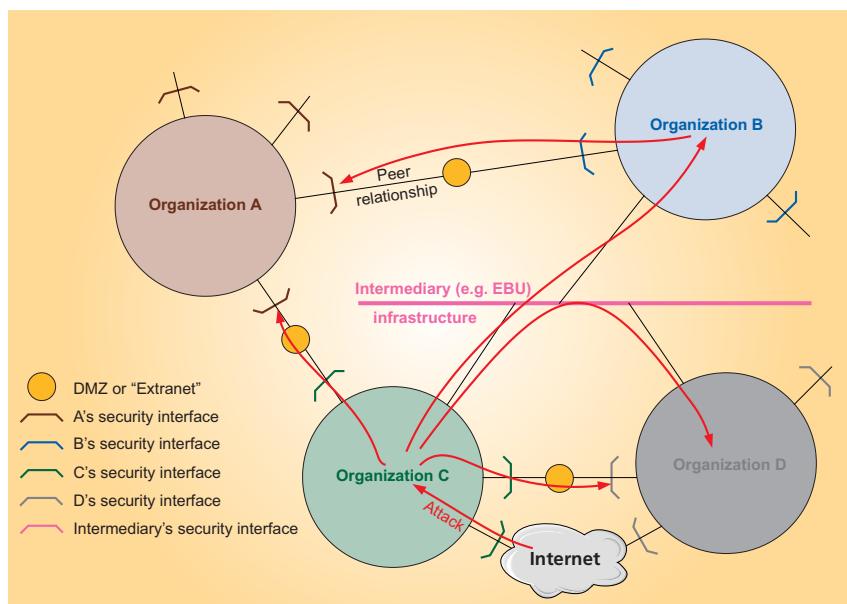


Figure 5
A lack of security interfaces allows the attack to hop-through and spread

5.1. *Broadcaster to/from broadcaster in a peer relationship*

Broadcaster (or production company) **A** may choose to form a direct and equal relationship with broadcaster (or production company) **B**. This relationship would enable the two organizations to exchange material and information. This is demonstrated in *Figs. 3 to 5* between organization **A** and organization **B**.

From a security perspective, the two organizations will need to agree common standards and policies regarding their interactions. For a large organization, this leads to a scaling issue. If the broadcaster interacts with n other organizations, they will need to agree $(n^2 - n)/2$ relationships, which will all need to be different. This can be a burden, which is why most companies adopt a standard “shell” security model. On the other hand it also lends a strong argument towards the creation of an inter-broadcaster exchange service.

5.2. *Broadcaster to/from broadcaster in a customer-supplier relationship*

One broadcaster (or Production Company) may routinely supply information or material to another as part of a contractual agreement. In extreme cases, the supplier may be 100% dependent on their customer, which would have an impact on how the relationship is realized.

The security aspects are similar to those described in *Section 5.1*. However, a relationship where one supplier is wholly working for one customer means that the simplest approach is for the supplier to wholly adopt the customer's policies and processes (effectively being invited inside the customer's security shell).

5.3. *Broadcaster to/from an intermediary organization*

An organization might operate as a material- or information-exchange service. In purest terms, the organization will not create their own material, but will instead facilitate the movement of material and information between any two, or more, of its members. This is illustrated in *Figs. 3 to 5* between organizations **B**, **C** and **D** and the intermediary infrastructure.

From the intermediary organization's security perspective, this is similar to the case described in *Section 5.1*. – with n members, they will need to establish $(n^2 - n)/2$ relationships. However from the member's point of view, they will be able to have a relationship with all the other members but will only need to establish one security relationship (with the intermediary organization). Extra care will be needed, because a security flaw affecting the intermediary could potentially affect all the members due to hop-through attacks.

5.4. *Combinations of the above relationships*

Many organizations will find themselves operating a combination of the above relationships. For example, a broadcaster may receive contributions from a number of production houses, some of which exclusively supply that broadcaster. The same broadcaster may also supply material that it has produced (along with raw material needed for other productions) to a different broadcaster. This same broadcaster might operate a co-production with a broadcaster based on the other side of the globe. And again, the same broadcaster might also submit material to, and receive material from, an intermediary organization such as the EBU.

Such a combination of relationships will mean that a typical medium-sized broadcaster could be trying to manage a large number of security boundaries. For cost-efficiency purposes, most will adopt a layered shell approach and agree variations (or build Extranets) with other individual broadcasters or intermediaries as needed.

6. The entities that flow between two broadcasters

Fig. 6 is an attempt to classify the sorts of “flows” that might be needed between any two broadcasters.

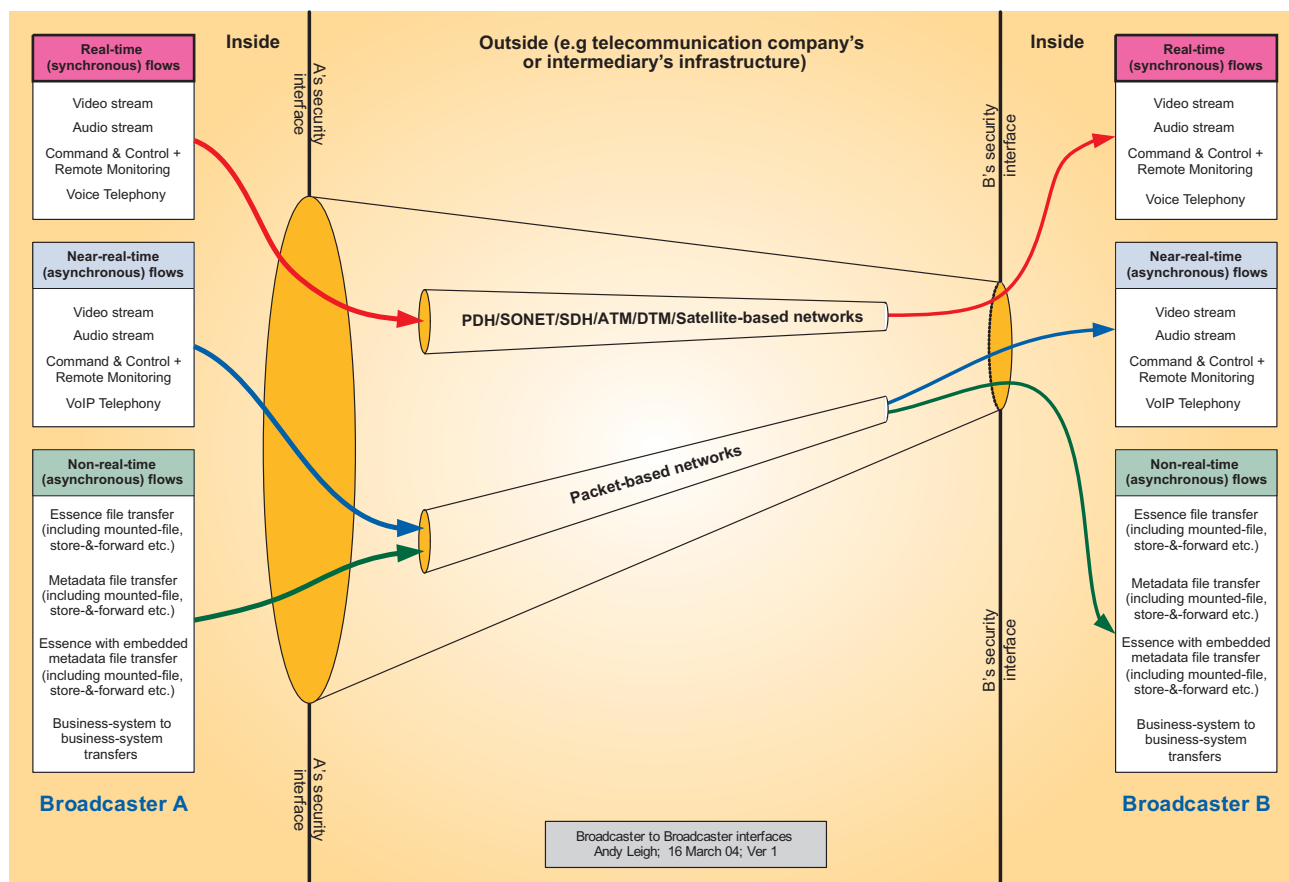


Figure 6
The flows between two broadcasters

6.1. Protecting an audio or video port

Traditionally, no attempt has been made to ensure that the interfaces between broadcasters are built securely. To understand why this is not a problem, we need to consider an audio or video port connecting two such organizations together:

- **Principle 1 – Be simple and easy to understand. [Pass or Fail? = Pass]**
Although broadcast systems (analogue and digital) are complex in themselves, the interfaces have always been relatively simple with extremely limited in-band signalling and self-routing capability.
- **Principle 2 – By default, block all access to it and through it unless that access is expressly permitted. [Pass or Fail? = Pass]**
An audio or video port will do very little if a faulty signal is applied to it. It certainly won't allow access into the device or access through it.
- **Principle 3 – Ensure that access to it, and transactions through it, are authenticated (i.e. performed by people or systems that have uniquely identified themselves and authenticated themselves, e.g. through a password certificate or key). [Pass or Fail? = Pass]**
See Principle 2 above.
- **Principle 4 – Be based on an open design, not secured through obscurity. [Pass or Fail? = Pass]**
Audio and video port designs are public knowledge.

- **Principle 5 – Have a completely locked-down configuration that can only be altered by both organizations with mutual consent (not by one of them alone). [Pass or Fail? = Pass]**
Access via the port to the configuration engine is not possible and since the system is physically secured, access by the other party is not possible.
- **Principle 6 – Restrict access to only those agreed facilities that are absolutely necessary but to no others. [Pass or Fail? = Pass]**
Hopping through an audio or video switch to gain access to other systems is not possible.
- **Principle 7 – Only utilise the bare minimum of shared apparatus between the two organizations. [Pass or Fail? = Pass]**
There is only a cable carrying a narrowly defined communications protocol.
- **Principle 8 – Be easy to use or transparent to the people and systems that (i) need to use it, (ii) are authenticated and (iii) can gain access to the objects they need. [Pass or Fail? = Pass]**
Beyond the constraints associated with needing to secure frame rooms and buildings, there are no psychological implications due to securing the technology.

So for decades, broadcasters have unintentionally been using and building interfaces that meet all eight of Saltzer's & Schroeder's principles. Added to the cost of entry for an attacker, it is no wonder that broadcast systems have never had to take security into account until now.

6.2. *Protecting an interface built from a computer with two network cards*

If we instead apply the same rules to a black box (see Section 4.) based on an IP-to-IP gateway built from a PC running an unmodified Operating System and with two network cards (*similar to Fig. 1*), then we encounter a completely different set of results when we apply the principles:

- **Principle 1 – Be simple and easy to understand. [Pass or Fail? = Fail]**
The computer is based on a multi-user, multi-tasking Operating System. There are hundreds or thousands of libraries linked in the kernel. Many services are automatically started on boot-up. A lot of software is installed and runs by default. Multiple TCP and UDP ports are automatically opened. Routing might be enabled. There may be applications running and bound to a TCP/UDP port.
- **Principle 2 – By default, block all access to it and through it unless that access is expressly permitted. [Pass or Fail? = Fail]**
It is possible to ensure the facilities on the box and through the box are locked down unless permission is given to access them. However, if there is a means for someone to pretend to be the administrator, then local access will be possible. Also, if the box allows a connection through to another computer, there is a risk that this computer and all others will be vulnerable.
- **Principle 3 – Ensure that access to it, and transactions through it are authenticated (i.e. performed by people or systems that have uniquely identified themselves and authenticated themselves, e.g. through a password certificate or key). [Pass or Fail? = Fail]**
Where the numbers of people are small or they all work for one organization, it is possible to ensure that facilities on the device and through the device are only accessible by authenticated people. Allowing multiple people or organizations to have access requires a complicated ID and password system. Standards for a solution (known as "Federated Identity") are still being developed. Also, if there are any software flaws in any of the programs or the Operating System that would allow access without a password, then control would be lost.
- **Principle 4 – Be based on an open design, not secured through obscurity. [Pass or Fail? = Pass]**
Most gateway-type solutions are based on open designs.
- **Principle 5 – Have a completely locked-down configuration that can only be altered by both organizations with mutual consent (not by one of them alone). [Pass or Fail? = Fail]**

It is very hard to build a PC-based solutions that allows such granularity of control. If one person knows the administrator password, they will be able to make any changes without the consent of the other party. It is for this reason that most solutions use two back-to-back gateways with a “DMZ” between.

○ **Principle 6 – Restrict access to only those agreed facilities that are absolutely necessary but to no others. [Pass or Fail? = Fail]**

A PC gateway runs a large number of services. The networks beyond the PC also run a very large number of functions. Building rules that restrict access is a very complicated task and in some cases, a suitable set of restriction rules can never be achieved.

○ **Principle 7 – Only utilise the bare minimum of shared apparatus between the two organizations. [Pass or Fail? = Fail]**

If the gateway is based on a single PC, then the solution suffers from the complexity described in Principle 1 above, but the situation is worsened because both organizations share that complexity. However, if the solution utilises two back-to-back devices with a DMZ between them, then the complexity of the PC is no longer shared. Instead, the complexity of the IP communication link is shared. Any IP path has, by default, more than 65,000 TCP connections and more than 65,000 UDP connections potentially live at any time. In practice the only ones of concern are those bound to software in either PC. Consequently, all services and all applications on the PC need to be reviewed and, if possible, removed.

○ **Principle 8 – Be easy to use or transparent to the people and systems that (i) need to use it, (ii) are authenticated and (iii) can gain access to the objects they need. [Pass or Fail? = Partial Pass]**

It is possible to achieve an element of transparency but, where people are used to connections just working, any restriction always causes friction and people do put effort into subverting the gateway just to “get their job done”. At the same time, gateways of this nature often impose performance strictures and this can lead to problems.

From this analysis, we conclude that it is not possible to build a trustworthy interface between two organizations using standard PC-type products and software out of the box. We have also shown that such a file or stream-transfer interface cannot safely be used as an alternative for a dedicated audio or video port. It is possible to minimise the risk, but excellent network, OS and software practices must be used during the design and deployment, and must be maintained for the life of the system.

7. Looking at security from a “layers” perspective

The OSI 7-layer model is frequently used inappropriately outside of its simple network architecture space. However, it has become a familiar descriptive methodology for demonstrating a hierarchical service-based approach.

Fig. 7 shows some of the security elements that need to be considered and how they are layered relative to each other. Security cannot be achieved by simply performing one task – such as encrypting data or installing a filtering firewall. In almost all cases, a mixture of all of the solutions below is required to achieve a security level that meets the organization’s needs.

From the above, we can observe that only about two thirds of a secure facility can be attributed to the deployment of technology. Also, each technology has a relatively narrow impact, but policies, processes, operations etc. affect every aspect and every layer. Interestingly, many experts consider that the critical success factors to building secure facilities are: (i) clearly defined security (and acceptable-use) policies, agreed and signed at the highest level; (ii) clearly defined and up-to-date security processes and standards including change-management and (iii) operational staff that are well trained and experienced in dealing with security processes.

We can argue that everything below layer 7 can be considered as **Networking** (blocked in green in *Fig. 7*) and everything above layer 7 can be considered as **Services & Applications** (blocked in

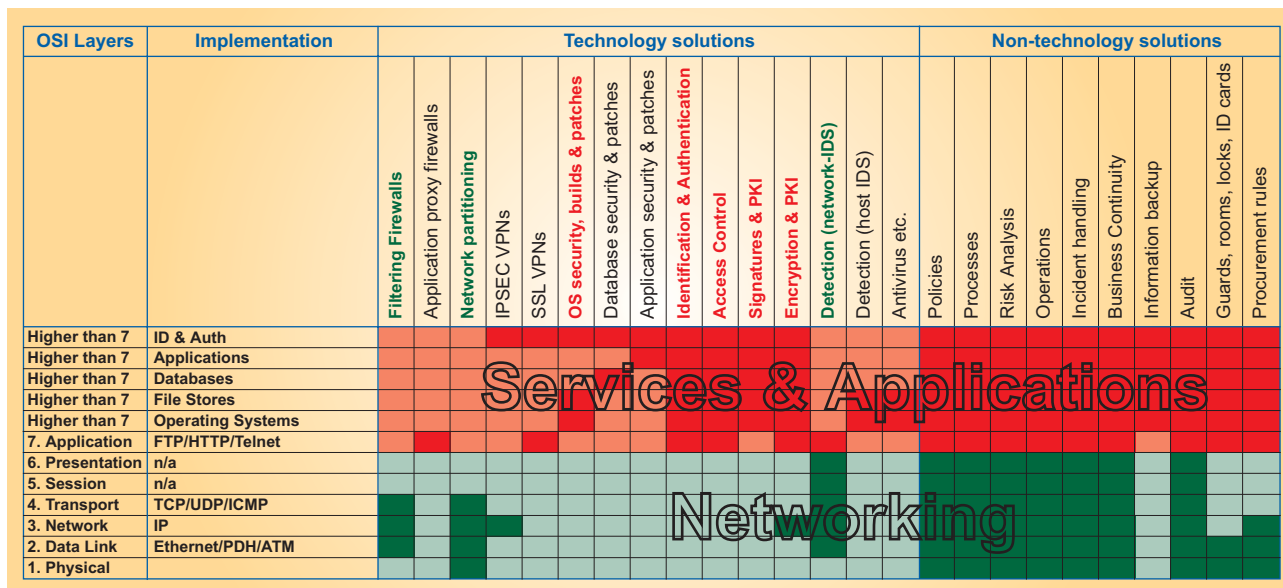


Figure 7
Layers and applicable tools

red). There are more red technology blocks than green technology blocks, which seems to suggest that secure network components do not have as great a part to play as secure services and applications. The *network technology* columns that have more than one solid green block are: Filtering Firewalls; Network partitioning and Detection (network-IDS). By combining the network policies and the network technologies, we can conclude that a network designed to enable secure working:

- utilises separate networks for separate tasks joined by filtering firewalls;
- is installed in secure rooms and cable ducts;
- operates encrypted links delivered by technologies such as IPSEC VPNs;
- has had a risk analysis performed on network designs, developments and changes;
- is designed around sound security policies and processes, and is maintained and operated securely;
- is monitored for intrusions, which are acted upon;
- has a highly resilient architecture;
- is regularly audited to ensure it remains secure.

But such a network will *not* make an organization secure. This is because the services and applications must also be accounted for. From a technology perspective, the columns which have the greatest number of red blocks are: Identification & Authentication; Access Control; Signatures & PKI; Encryption & PKI, and OS security, builds and patches. By combining the application and service *policies* and the application and service *technologies*, we can conclude that a system designed to enable secure working:

- is based on securely-designed OS builds that are regularly patched and upgraded;
- is based on securely built databases that are regularly patched and upgraded;
- is based on securely built applications that are regularly patched and upgraded;
- identifies and authenticates any user or system that needs access;
- refuses access to anyone who is not authenticated;
- signs or encrypts files, streams and data stores where necessary;
- is installed in secure frame rooms;
- protects against malware;
- has had a risk analysis performed on the Operating System and database, and application designs, developments and changes;
- is designed around sound security policies and processes, and is maintained and operated securely;

- is monitored for anomalous behaviour, which is acted upon;
- has a highly resilient architecture;
- is regularly audited to ensure it remains secure.

Many of the current generation of IT-based broadcast systems do not meet these criteria. They could be built to do so, in many cases with almost no impact on development or operational cost. However, because the broadcast industry is not yet aware that the facilities it procures need to have security built into the development and operational lifecycles, no current or future systems will be secure. The potential impact is that inter-broadcaster communications will not be able to move over to new flexible ways of inter-working without taking unacceptable risks with their most important technical and content assets. To reap the long-term improved efficiencies that many consider to be essential for the continued growth of the industry, some planning and investment is needed now.

It is therefore incumbent on broadcasters and production companies to build, procure and operate systems and networks that are intrinsically secure. It is also incumbent on suppliers and integrators to design and build systems and networks that are intrinsically secure and that can be consistently maintained as secure.

8. Deriving inter-organizational solutions from these flows and layers

If two organizations need to work together, they should first consider the flows that will need to move between the two entities (*see Section 6.*). This process must be very specific on exactly what objects need to be shared and the timeliness of that sharing. It will be extremely difficult to supply a trustworthy solution to a general requirement such as “we need a network link between us and them”.

The organizations should then consider who or what needs access to the objects that will be shared. It will be extremely difficult to supply a secure solution to a general requirement such as “we need to let their staff see our schedules”. Named individuals or systems should be nominated and a means found to ensure that both organizations can satisfy themselves that these individuals and systems have been uniquely identified and authenticated. There needs to be clarity on whether all the specific objects should be available for any of the nominated individuals or systems to **edit**, or whether this only applies to a subset, and that **read** access will be sufficient for most.

Now that both organizations understand the flows, they can turn to the layers of technology that should be applied (*see Section 7.*). The principles outlined in *Section 2.* should also be applied to any technology utilised to enable the inter-organizational flows. The solution should start with policies, processes and standards and should also ensure that the infrastructure is built and managed by trained and competent staff.



After graduating with a degree in Electronics from the University of Manchester Institute of Science and Technology (UMIST, UK) in 1985, **Andy Leigh** joined the BBC where he started as a broadcast engineer working for BBC Radio. After about five years, he took on a new role in computing which also gave him the responsibility for operating BBC Radio's newly installed Ethernet data network. Subsequent changes saw him take on the operation of all the BBC's data networks where he also became involved in their design and development.

In the mid 1990s, Mr Leigh began to work for the Strategic Network Development department where he concentrated on the development of protocols and the integration of broadcast systems with more traditional IT technologies. His work on network strategy drew him into consultancy work with the BBC's Information Security team and, in 2002, he became the BBC's Information Security Strategist. He is a founding member of the Jericho Forum and chairs the EBU N/Security working group.

Conclusions

Designing and agreeing the security relationships between broadcasters and production companies is a complex process. Care needs to be taken over the definition of the various relationships which might be pursued. A number of different entities can be seen to flow in different relationships, and a wide range of technologies and processes can be utilised to secure these flows.

The flow-to-layer mapping process (*from Section 8.*) could form the basis for defining a set of standards that all broadcasters and intermediaries agree to adopt. Details of the best approaches and technology could then be agreed and applied to any specific relationship that is required.

References

- [1] Jerome Saltzer and Michael Schroeder (MIT): **The Protection of Information in Computer Systems**
Proceedings of the IEEE, 63(9), Sept. 1975. Available via <http://www.ieee.org/portal/site>
 - [2] S. Kent and R. Atkinson (editors): **Security Architecture for the Internet Protocol**
RFC2401
<http://rfc.net/rfc2401.html>
 - [3] John Postel (editor): **Transmission Control Protocol**
RFC793.
<http://rfc.net/rfc793.html>
 - [4] John Postel (editor): **User Datagram Protocol**
RFC768
<http://rfc.net/rfc768.html>
 - [5] Bruce Schneier: **Secrets & Lies – Digital Security in a Networked World**
John Wiley & Sons, 2000.
<http://www.wiley.com/WileyCDA/WileyTitle/productCd-0471453803.html>
 - [6] CERT Coordination Center: http://www.cert.org/stats/cert_stats.html
-