An Evaluation Framework for the Analysis of Covert Channels in the TCP/IP protocol suite

D Llamas, A Miller and C Allison
School of Computer Science, University of St Andrews, St Andrews KY16 9SX, Scotland, UK
Tel: +44 (0) 133 4463253, Fax: +44 (0) 133 4463278, {david,alanr,ca}@dcs.st-and.ac.uk

Abstract: Information hiding techniques can be used by criminals and terrorists to communicate over covert channels within the TCP/IP protocol suite and can be used to overcome firewalls and most other forms of network intrusion detection and prevention systems. In this work we describe the covert channel concept and weaknesses in the five layered TCP/IP layered model. We then present an evaluation framework for the analysis of covert channels and illustrate this with an example featuring the heavily used IPv4 datagram header.

Keywords: Information Hiding, Covert Channels, Evaluation Framework, TCP/IP

1. Introduction

The covert channel concept was introduced in 1973 (Lampson 1973). A covert channel is a communication channel that allows two cooperating processes to transfer information in a manner that violates the system's security policy (Berg 1998). It is a means of communication that is not part of the original design of the system. It can be used to transfer information to a process or user that, a priori, would not be authorised access to that information. Covert channels only exist in systems with multilevel security (Proctor and Neumann 1992), which contain and manage information with different sensitivity levels. A covert channel allows different users access to the same information, at the same time, but from different points-of-view.

With the expansion of the Internet into every corner of the Globe, the protocols which underpin it provide a potentially ubiquitous environment for covert channels, which has received little attention. The contributions of this paper are threefold. An evaluation of covert channels in relation to each layer of the Internet architecture is presented. A classification of network protocol covert channels, which facilitates both systematic categorisation and focussing on areas of perceived threat is discussed. An evaluation framework which can be applied to the specific potential covert channels identified is outlined. A covert channel which we have implemented in IP is then outlined and the evaluation framework applied to it.

2. The TCP/IP Layered Model

TCP/IP is a set of network protocols developed for the Internet from the 1970s. As a protocol suite based on layers, TCP/IP has a number of weaknesses that allow an attacker to leverage techniques in the form of covert channels to surreptitiously pass data in otherwise ordinary packets (Llamas 2004). Many techniques are based on encoding data in the protocol header (daemon9, AKA et al. 1996; Handel and Sandford 1996; Route 1996; Rowland 1996; Abad 2001; Ahsan 2002; Giffin, Greenstadt et al. 2002).

The appropriateness of protocol layers for covert channels are evaluated with respect to three criteria, which we name, *technical difficulty*, *generality* and *reachability*.

- **Technical Difficulty:** What are the technical barriers required to establish and read a covert channel? Does it require special hardware, alteration of the operating

system or low level programming, programming in the application space or simple system configuration?

- **Generality:** Once the technical barriers for a covert channel have been overcome how widely can they be applied? Is all Internet traffic susceptible or only some subset thereof?
- **Reachability:** If a covert channel is established how far can it reach through the Internet? For example, is it likely to be confined within an institution or be on a global scale?

In this section a high level description of each layer of the TCP/IP protocol stack is presented. The characteristics and potential for covert channels of each layer is briefly discussed (Girling 1987; Handel and Sandford 1996). In the protocol stack each layer is implemented by one or more protocols (horizontal) and one or more interfaces (vertical), as outlined in Figure 1.

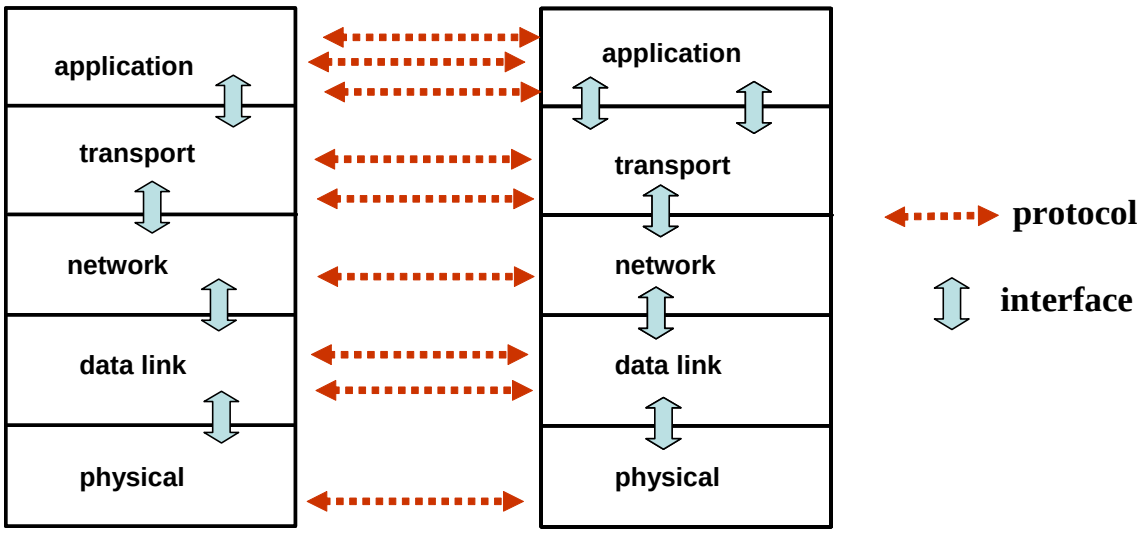

Figure 1: The TCP/IP 5-Layer Model

Physical layer

The Physical Layer is concerned with the physical media used to directly connect adjacent nodes. The physical layer's job is to transfer signals from one node to the next. The protocols in this layer depend on the actual transmission medium used.

There are a number of barriers that make this layer a difficult environment for the establishment of covert channels. There would be a significant overhead in designing the mechanisms for encoding and decoding the covert channel. This would probably require specialist hardware to be installed. What is more, the diversity of protocols and media mean that different technologies would have to be developed for each type of network, severely limiting generality. Furthermore, the reachability of the covert channel would be limited to adjacent nodes with some limited extension possible where repeaters are in use. Reachability would often be confined within an institution or some subset thereof.

Link layer

The link layer is responsible for the delivery of frames between adjacent nodes on a network. It provides delineators for the start and end of a frame, a link layer address, and usually some reliability bits which enable corruption to be detected. The link layer also includes the Media Access Control, (MAC) sub-layer, which allows contention between nodes competing for access to the medium to be resolved.
There are a number of technical barriers that would need to be overcome to establish covert channels at this level. Specialist hardware or low level device driver programming are likely to be necessary. Different types of network have different Link Layer protocols. For example, Ethernet, ATM and Token Ring. The reachability available to covert channels at this layer will often be limited to the networks within an institution as there are logical and physical limits on simply extending LANs such as Ethernet, and eventually a level-3 network router is used.

Two potential ways in which the link layer may be used for covert channels are given next (Handel and Sandford 1996). Firstly, the collision detection system (Carrier Sense Multiple Access / Collision Detection CSMA/CD) in the traditional Ethernet link layer can be modified to transmit hidden data by adjusting the collision control mechanism. Secondly, unused portions of the frame can be used to store covert data. Covert data can be stored in the buffer, beginning at the end and working toward the valid data. When the packet is transmitted, the entire buffer is exported, including the covert data.

Network layer

The network layer is responsible for routing datagrams from the source to the destination host. It provides network addressing and data routing. The IP protocol defines the addressing system and how intermediary nodes should treat datagrams. Each packet can be independently routed from source to destination. It provides a best effort service to higher layers. There are no guarantees that a datagram will be delivered, that a sequence of datagrams will be delivered in order, how long delivery will take or about the variation of delay that can be expected between datagrams. The IP protocol is fundamental to the Internet, every host attached to the Internet must have an implementation. In fact the Internet can be defined as the sum of hosts that are globally reachable through IP. The network layer also contains control and routing protocols.

Covert channels can be created solely in software at the Network layer, however this is likely to require some device driver programming or alteration of operating system code. Developing a single technology to establish covert channels in the IP layer, would yield a high degree of generality. Similarly IP level covert channels offer global reachability. The opportunity for discovery of IP level covert channels are enhanced by the simplicity of the IP header and options.

Analysis of the IP header shows the existence of bits that are either unused or optional. Consequently, fields from the IP header can be manipulated to store covert data. An extended analysis of using the IP Identification field as a covert channel is given later in this paper.

Transport layer

Transport protocols are the lowest true “End to End” protocols. They facilitate the delivery of data from the sending process on one computer to the receiving process on another computer. On the Internet they build upon the best effort service provided by IP. There are two main Internet transport protocols the Universal Datagram Protocol (UDP) and the Transmission Control Protocol (TCP). Both protocols contain port numbers, which enable multiple channels of communication to be distinguished between the same two hosts. TCP builds upon IP to provide a reliable, point-to-point, duplex, stream based virtual connection. TCP establishes and terminates connections between hosts and uses a system of acknowledgements and retransmissions (Howe 2003) . TCP also provides flow control and congestion control for the connection thereby ensuring the destination can keep up with the source and that network resources are fairly shared between multiple hosts. UDP provides an unreliable best effort service and acts as a placeholder which facilitates the development of higher level transport protocols such as the Real Time Protocol (RTP).

The location of transport protocol implementations within the operating system, make the technical difficulties in creating covert channels similar to the Internet layer. The creation of bespoke device drivers or the modification of operating system source code is likely to be necessary. It is variously estimated that over 80% all Internet traffic (Lam 2005) is carried by TCP therefore establishing a covert channel within TCP provides a high level of generality. The fact that TCP is an end-to-end protocol gives it global reachability. In addition, the TCP header is more complex than IP and has a significant array of options, which means there is more potential for the creation of such channels.

Two examples of TCP covert channels (Route 1996; Rowland 1996) are given below:

- **Initial Sequence Number Field:** The Initial Sequence Number field (ISN) of the TCP/IP protocol suite enables a client to establish a reliable protocol negotiation with a remote server. As part of the negotiation process for TCP/IP, several steps are taken in what is commonly called a "three way handshake”. For our purposes the sequence number field serves as a good medium for transmitting clandestine data because of its size (a 32 bit number). In this light, there are a number of possible methods to use. The simplest is to generate the sequence number from our actual ASCII character we wish to have encoded.

- **The TCP Acknowledge Sequence Number Field "Bounce":** This method relies upon basic spoofing of IP addresses to enable a sending machine to "bounce" a packet of information off of a remote site and have that site return the packet to the real destination address. This has the benefit of concealing the sender of the packet as it appears to come from the "bounce" host. This method could be used to set up an anonymous one-way communication network that would be difficult to detect especially if the bounce server is very busy. This method relies on the characteristic of TCP/IP where the destination server responds to an initial connect request (SYN packet) with a SYN/ACK packet containing the original initial sequence number plus one (ISN+1).

Application layer

This layer handles issues like network transparency, resource allocation and problem partitioning. The application layer is concerned with the user's view of the network, such as formatting electronic mail messages (Howe 2003). The Application Layer contains a diverse range of protocols for applications such as email, remote administration, World Wide Web (WWW) and Peer to Peer content distribution. An important application is DNS (Domain Name System) which translates host names into IP addresses.

The technical difficulties in creating covert channels are easiest to overcome at the application layer, limited or no programming skills may be required. The diversity of protocols make the generality of a technical solution limited, however most networked computers can be expected to support common protocols such HTTP and SNMP. Reachability will be global.

The application layer is nearest the user. Users create applications utilising system resources, including the network. Many of the classical steganographic approaches can be used at the application level. What would originally be done with a pencil and paper can be done here. For example, a covert messaging system can be devised using word substitution in an email system.

Summary

Opportunities for the creation of covert channels exist at all layers in the TCP/IP protocol stack.

The technical difficulties to overcome are highest at the physical and link layers and may require specialist hardware. At the Network and Transport Layers, access to and the ability to modify device drivers or operating system support code is required. At the application layer, special access privileges and programming skills are unlikely to be required. The higher the layer, the easier the access and the less skills required. Due to the nature of layered network models access at a lower layer can be used to introduce a covert channel at an arbitrary higher layer.

Internet protocol diversity can be visualised as a bell shaped structure. There is much diversity at low and high levels. There is only a single protocol at the network level and two protocols at the transport layer. The implication for generality is that establishing a covert channel at the network or transport layer will achieve the highest level of general coverage.

The level of reachability increases as one travels up the network stack. Covert channels at the physical and link layers will achieve limited reachability. Those at the network layer and above will potentially achieve global reachability.

The above considerations lead us to focus on covert channels in the network layers and above. Reasons of tractability lead us to focus on the Network and Transport layers in the rest of this paper.

3. Network Covert Channels Classification

We propose a two type classification of Covert Channels in computer networks. At the first type we identify covert channels that rely upon the manipulation of protocol

header bits to communicate information and channels that modify the behaviour of flows to communicate information. Examples of modifying the behaviour of flows to establish covert channels may include, controlling the size of packets, the timing of packets and the destination of packets. The best effort service provided by TCP introduces noise into such channels and limits the bandwidths that are achievable.

TCP/IP is a set of protocols based on a layered model. The content of each level is wrapped by a header, which contains the information needed for the particular purpose of that protocol. Header fields may contain some bits that are redundant and other bits that can be overloaded to allow them to be utilised in the creation of covert channels. For the rest of this section we focus on the manipulation of bits in the header.

At the second type of classification we identify two criteria. Whether the covert channel can be created without violating the legal framework of the protocol and whether the channel can be created within the constraints of actual implementations of the protocol. This possibility arises because protocol implementations and their specifications are not identical. Consequently, covert channels could follow the legal definition of the field to be manipulated but might not work at the implementation level, or they could not strictly follow the legal definition but work at the implementation level. This in turn gives four types of covert channel.

- Legal and implementation supported: Will be difficult to detect and have a high utility
- Legal and are not implementation supported: Will be difficult to detect but have limited utility
- Illegal but are implementation supported: Will be easier to detect and have a high utility
- Neither legal nor implementation supported. Will be easier to detect and have limited utility

The first phase of our work is subjecting the TCP and IP protocol headers to a rigorous analysis to determine the opportunity for establishing covert channels within the legal framework of these protocols, then how well a particular covert channel will work within the selected field or fields in that protocol.

In the Internet there will be multiple implementations of the same protocol running simultaneously. Covert channels aim to use protocols in ways that they were not intended. It is possible for a covert channel to interact unpredictably with different implementations of protocols. Consequently it is important to measure to what extent a covert channel can work with the deployed protocol implementations base. This can be achieved by conducting live Internet experiments against the base of actually deployed protocols.

The combination of analysis of packet headers and experimental work allows the identification of a number of potential covert channels. These will share the properties of being hard to detect and having a high potential utility.

4. Evaluation Framework

So far in this paper the focus has been on the general characteristics on protocol layers and how they relate to covert channels and a broad categorisation of covert channels. In this section an evaluation framework is outlined, which allows a more detailed analysis of specific covert channels to be undertaken. Three categories are addressed. The first concerns the generation of covert channels. The second concerns how covert channels can be detected and the third concerns the ease with which covert channels can be disrupted. An example is then given, featuring the IPv4 datagram header.

Generation of channels

If it is known how easy it is to generate a covert channel and the technical steps that are required, then precautionary measures can be taken to limit the opportunity for their creation. The following issues will be addressed when considering the generation of channels:

- Which aspects of these protocols lend themselves to this purpose?
- What level of access to normal computing network interfaces is required?
- What technical competence and skills are required to effect this exploitation?
- What bandwidths can be obtained?

Detection of channels

If a covert channel is detected it ceases to be a covert channel. However, there is a cost which is attached to the detection of such channels. In this part of the framework we address the cost of detecting covert channels.

- How can packet monitoring, filtering and analysis detect such channels?
- What levels of performance are required by a monitor and analyser?
- Up to what rate can this be done in real time?
- Where should monitoring points be located?

Disruption of channels

The particular characteristics of covert channels in digital network protocols, mean that effective countermeasures may be more oriented to prevention than detection. It may be possible to put measures in place which disrupt the performance of known or unknown covert channels. The most advanced Intrusion Prevention Systems (IPS) nowadays do not currently include this kind of security technique (Franklin and Wiens 2005). Two questions are relevant here:

- Is it possible to prevent covert channels by disrupting features of a protocol?
- Is it possible to disrupt transparently, so that the overt data flows and functions associated with the protocol are not adversely affected?

Example: Manipulation of the IP Identification Field

The identification field of the IP protocol (see Figure 2) helps with re-assembly of packet data by remote routers and host systems. Its purpose is to give a unique value to packets so if fragmentation occurs along a route, they can be accurately re-assembled. The encoding method described here simply replaces the IP identification

field with the numerical ASCII representation of the character to be encoded. This allows for easy transmission to a remote host which simply reads the IP identification field and translates the encoded ASCII value to its printable counterpart.

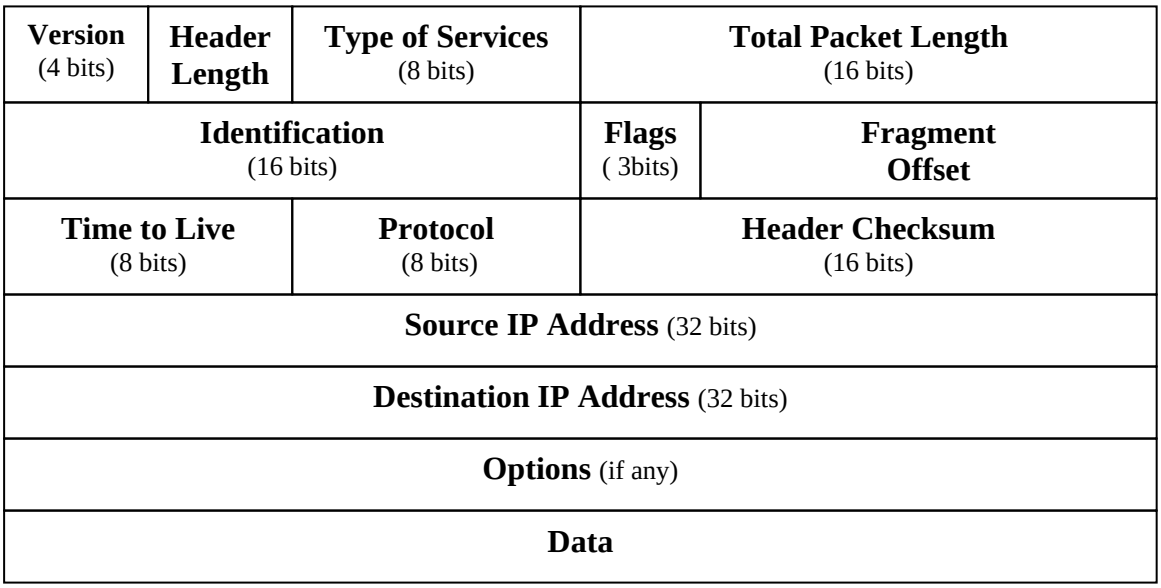

Figure 2 - IP header v4

This value is generated at the origin, which means that the *identification* field can be used to store a specific value to be passed to the other end. Using this method it is possible to pass data between hosts in any IP packet. These packets can contain no actual data, or can contain data designed to look innocent. These packets can also contain forged source and destination IP addresses as well as forged source and destination ports. This can be useful for tunnelling information past some types of packet filters as well as simply anonymising communications between sender and receiver.

• **Protocol**: IP • **Field**: Identification • **Type**: Numeric • **Length**: 16 bits • **Values**: From 0 to 65,535	• **Function**: An identifying value assigned by the sender to aid in assembling the fragments of a datagram (DARPA 1981). • **Weakness**: Number assigned by the sender with non required sequence.

A value from 0 to 65,535 can be given by the sender in each IP packet. Each of these values can have a meaning. For instance, some of them can be associated to different alphabets, others can be related to numbers, composed messages, times, places, special symbols and so on. As it has been explained above, the identification field on the IPv4 header is suitable for a covert channel to be created. We have implemented this covert channel and demonstrated its use and robustness against interference from conventional firewalls. Next we apply the evaluation framework proposed in this paper.

Generation

Which aspects of these protocols lend themselves to this purpose?

The ubiquitous nature of IP and the global reach of the protocol lend itself to use for covert channels generally. The existence of an IP identification field and the widespread use of the “do not fragment flag” being set allows this field to be reused as a covert channel. Strictly the legal definition of the protocol is broken as it is no longer guaranteed that all IP in flight IP packets can be uniquely distinguished. However, in practice the probability of a clash is low.

What level of access to normal computing network interfaces is required?
To set up the system for the generation of covert channels system administrator privileges are required so that bespoke device drivers may be installed or the operating system source code modified. Once installation had occurred only normal user level privileges are required.

What technical competence and skills are required to effect this exploitation?
Deep knowledge of the TCP/IP protocol suite and the IP protocol in particular are required in order to design, write and deploy the required software. To use the system once such installation has occurred the skills of a proficient user should suffice.

What bandwidths can be obtained?
The identification field in the IPv4 header is defined as a 16 bit number. A bandwidth of 16 bits per packet can therefore be achieved. This approximates to 1/700th of the bandwidth that an Ethernet flow would achieve as Ethernet frames usually contain 1460 bytes of application level data.

Detection
How can packet monitoring, filtering and analysis detect such channels?
In this particular case of the identification field, the covert channel is legal, but because the id number has to be unique during the life of the packet, the main implementations are incrementing one each time a packet is generated. For this reason, a way of detecting a potential covert channel in that field could be the observation that the identification number is not incremented by one. This can be achieved through automated passive monitoring of packets, combined with some simple on line analysis software.

What levels of performance are required by a monitor and analyser?
A standard Unix workstation with an off-the-shelf high performance network card would be able to undertake passive monitoring and detection of this covert channel.

Up to what rate can this be done in real time?
It can be done in real-time up to 100 mbit/s. At bandwidths of 1000 mbit/s per second high performance monitoring cards with specialised onboard packet filtering capabilities would be required.

Where should monitoring points be located?
The most suitable locations for maximising the detection of this particular covert channel would be at the Ingress and Egress of institutional networks. At this point covert channels which attempted to communicate outside of the institution could be detected.

Disruption

Is it possible to prevent covert channels without disrupting features of a protocol?

The covert channel presented as the example, based on the identification field of the IP v4 header, could be disrupted by overwriting the field with a new number in line with a standard implementation. Examples of disruption techniques may include rewriting all IPv4 identification fields of network packets with new numbers. This operation could be done at the Egress point of the network thereby ensuring that sensitive information is not exported.

Is it possible to disrupt covert channels transparently, so that the overt data flows and functions associated with the protocol are not adversely affected?

This way of disruption will avoid overt data flow and functions associated with the protocol being affected.

5. Conclusion

This paper has presented an analysis of the potential for covert channels in each layer of the TCP/IP protocol stack. The conclusion of the analysis was that the Network and Transport layers are the most vulnerable to the creation of Covert Channels. This paper has focused on issues relevant to the IPv4 and TCP protocols.

A two level classification of covert channels has been presented. Firstly, covert channels can be created by changing the behaviour of protocols. For example changing bandwidth usage, packet size, delay between packets or by modifying the contents of packet headers. The modification of the behaviour of these channels is not a desired route to look into due to the inefficiency and noise Secondly, for content based channels, we identify two important distinguishing characteristics. Whether the channel is within or violates the legal definition of the protocol and whether the channel would be supported by the deployed implementations of protocols. We observe that the second criteria can only be determined through live Internet experiments. We conclude that covert channels which are legal and are supported by deployed implementations are the most dangerous as they combine the characteristics of being potentially hard to detect and offering the user practical utility.

An evaluation framework, which can be applied to a specific covert channel to assess its practicality of generation and difficulty of detection or disruption was then proposed. The evaluation framework was illustrated by applying it to an IP layer covert channel.

6. Acknowledgement

This work has been performed in the framework of the PhD research programme of David Llamas at the University of St Andrews, which is supported by the Defence Science and Technology Laboratory - DSTL – a part of the Ministry of Defence of the United Kingdom.

7. References

Abad, C. (2001). IP Checksum Covert Channels and Selected Hash Collision. USA, University of California.

Ahsan, K. (2002). Covert Channel Analysis and Data Hiding in TCP/IP. Canada, University of Toronto.

Berg, S. (1998). Glossary of Computer Security Terms. USA, National Computer Security Center.

daemon9, AKA, et al. (1996). Project Loki. Volume Seven, Issue Forty-Nine. USA, Phrack Magazine.

DARPA (1981). RFC 791 - Internet Protocol. USA, Defense Advanced Research Projects Agency.

Franklin, C. and J. Wiens (2005). Intrusion-Protection Systems. The Great IPS Test. USA, Network Computing, CMP Media LLC.

Giffin, J., R. Greenstadt, et al. (2002). Covert Messaging Through TCP Timestamps. USA, Massachusetts Institute of Technology - MIT.

Girling, C. G. (1987). Covert Channels in LAN`s. USA, IEEE Transactions on Software Engineering.

Handel, T. G. and M. T. Sandford (1996). Hiding Data in the OSI Network Model. USA, Weapon Design Technology Group, Los Alamos National Laboratory.

Howe, D. (2003). FOLDOC Computer Dictionary. USA, Webnox Corp.

Lam, S. (2005). Back to the Future Part 4: The Internet. USA, ACM SIGCOMM Computer Communication Review.

Lampson, B. W. (1973). A Note on the Confinement Problem. USA, Xerox Palo Alto Research Center.

Llamas, D. (2004). Covert channel analysis and data hiding in the TCP/IP protocol suite. Honours Project Thesis. UK, Napier University.

Proctor, N. E. and P. G. Neumann (1992). Architectural implications of Covert Channels. USA, Computer Science Lab, SRI International.

Route (1996). Project Loki: ICMP Tunnelling. USA, Phrack Magazine.

Rowland, C. H. (1996). Covert channels in the TCP/IP protocol suite. USA, Vol.2 No.5 - First Monday Magazine.

8. Appendix

Authors biographies

David Llamas is a PhD student in the School of Computer Science at the University of St Andrews. His research areas of interest are information hiding systems: covert channels & steganography, architectures and protocols for public communication networks, evolutionary computing: learning classifier systems and aerial & underwater acoustic networks. He is the administrator of the specialised websites http://www.steganography.org and http://www.acousticnetworks.org

Web: **http://www.dcs.st-andrews.ac.uk/~david**
Email: **david@dcs.st-and.ac.uk**

Dr. Colin Allison is a Senior Lecturer in the School of Computer Science at the University of St Andrews, and leads the Distributed Systems and Networks research group. His research areas include distributed learning environments and holistic Quality of Service.

Web: **http://www.dcs.st-andrews.ac.uk/~colin**
Email: **ca@dcs.st-and.ac.uk**

Dr. Alan Miller is a Lecturer in the School of Computer Science at the University of St Andrews. His research areas of interest are Internet systems research: traffic measurement and monitoring, TCP congestion control, user perceptions of network Quality of Service, constructivist adaptive E-learning and social & technological factors in the Internet's development.

Email: **alanr@dcs.st-and.ac.uk**