

Lecture Notes in Computer Science

1796

Bruce Christianson Bruno Crispo
James A. Malcolm Michael Roe (Eds.)

Security Protocols

7th International Workshop
Cambridge, UK, April 1999
Proceedings



Springer

Bruce Christianson Bruno Crispo
James A. Malcolm Michael Roe (Eds.)

Security Protocols

7th International Workshop
Cambridge, UK, April 19-21, 1999
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany
Juris Hartmanis, Cornell University, NY, USA
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Bruce Christianson
James A. Malcolm
University of Hatfield/Herfordshire
Computer Science Department
Hatfield AL10 9AB, England
E-mail: {B.Christianson/J.A.Malcolm}@herts.ac.uk

Bruno Crispo
CSELT SpA - Gruppo Telecom Italia
Via Reiss Romoli 274, 10100 Torino, Italy
E-mail: bruno.crispo@cl.cam.ac.uk

Michael Roe
Microsoft Research
St. George House, 1 Guildhall Street
Cambridge CB2 3NH, England
E-mail: mroe@microsoft.com

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Security protocols : 7th international workshop, Cambridge, UK, April
95 - 21, 1999 ; proceedings / Bruce Christianson ... (ed.). - Berlin ;
Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ;
Singapore ; Tokyo : Springer, 2000
(Lecture notes in computer science ; Vol. 1796)
ISBN 3-540-67381-4

CR Subject Classification (1991): E.3, F.2.1-2, C.2, K.6.5, J.1

ISSN 0302-9743

ISBN 3-540-67381-4 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag is a company in the BertelsmannSpringer publishing group.
© Springer-Verlag Berlin Heidelberg 2000

Typesetting: Camera-ready by author, data conversion by DA-TeX Gerd Blumenstein
Printed on acid-free paper SPIN 10720107 06/3142 5 4 3 2 1 0

Preface

Another year, another workshop. Here are the proceedings of the seventh Cambridge International Workshop on Security Protocols. All very well, you may think, but can there really still be anything genuinely new to say? Is it not just the same old things a tiny bit better?

Well, perhaps surprisingly, this year we discovered some radically new things beginning to happen. The reasons in retrospect are not far to seek: advances in technology, changes in the system context, and new types of consumer devices and applications have combined to expose new security requirements. This has led not only to new protocols and models, but also to known protocols being deployed in delicate new ways, with previous fragilities of watermarking and mutual authentication, for example, becoming desirable features. At the workshop we identified several of these developments and began to map out some lines of enquiry.

This volume brings you a selection of deliberately disputatious position papers, followed by not-quite-verbatim transcripts of the discussions which they provoked. As always, our purpose in making these proceedings available to you is the hope that they will move your thinking in an unexpected direction. If you find your attention caught by something here, if it makes you pause to reflect, or to think “why, that is just *so* wrong”, then good. We’re waiting for your mail.

Thanks to Stewart Lee and the University of Cambridge Centre for Communications Systems Research, for acting as hosts for the workshop, and to Roger Needham and Microsoft Research Limited (Cambridge), for providing us with the use of their meeting room and coffee machine.

Thanks also to Dorian Addison of CCSR and to Angela Leeke and Margaret Nicell of MSRL for keeping us organized, and our especial gratitude to Lori Klimaszewska of the University of Cambridge Computing Service for her Promethean transcription of the audio tapes, which this year featured “echo kickback as well as the usual distorted sound”.

Finally, each of us takes full responsibility for the accuracy and completeness of the material contained in these proceedings. Errors and omissions, on the other hand, are the responsibility of the other three fellows.

February 2000

Michael Roe
Bruce Christianson
Bruno Crispo
James Malcolm

Introductory Remarks

Michael Roe: We always try and have a theme and then people submit papers. The usual thing that happens when you're running a conference is that people resolutely submit a paper on the research they're doing that year, regardless of whether it's got anything to do with the theme or not.

What I thought was going to be the theme — because I'd been deeply buried in it for the previous year — was entertainment industry protocols. We're seeing a great shift in the use of the Internet from the kind of educational, military, and industrial use of networking towards home users buying entertainment, and this causes you to completely redesign protocols. I've been particularly looking at copy protection and digital watermarking, and when you look at those kinds of protocols you discover they're just not like the usual authentication and crypto key exchange we've been trying to do for the last fifteen years. They're fundamentally different, even the kind of asymmetric properties you want in a digital watermark — that anybody can verify it and see that this is copyrighted data and nobody knows how to change it — you think, oh that looks like public key cryptography, but then you discover, no it's not public key cryptography, it's something that's similarly asymmetric but it's different.

And so I thought a theme for the workshop could be how these completely new application areas cause us to come up with protocols that are completely different from what we previously discussed. But from people's submissions this didn't look to be what everybody else was doing.

The second theme that Bruce suggested was auditability of protocols: What do we need in a protocol in order to audit it? What does it mean to have audit support in a protocol? Bruce, it was your idea . . .

Bruce Christianson: Well when something goes wrong, sometimes the question is, what actually happened? And then you work out what state you ought to be in. But I think we know that in the protocol world there's often not a fact of the matter about what actually happened. So we need instead to have some way of agreeing a story we're all going to stick to, about what we're prepared to say happened. Some kind of integrity property for restoring the state. It seems to me that several different pieces of work have each got one corner of that particular problem, and it might be interesting to discuss some of the relationships between those approaches that aren't usually discussed.

Michael Roe: So it's going to be a mixed bag this workshop. A general overall theme still might be the changing environment and the changing application areas, it's just things have changed so much and become so diverse that they exceed my ability to predict what the papers are going to be about.

Une Mise en Thème

an exchange of e-mail

stardate February 1999

From: B.Christianson@herts.ac.uk Fri 3 February 1999 14:45
To: E.S.Lee@ccsr.cam.ac.uk, m.roe@ccsr.cam.ac.uk
Subject: protocols workshop

there is some discontent with the present 'theme':

> one of the strengths of the earlier workshops was that they didn't
> concentrate upon a particular application area and so encouraged
> people with different interests to come together.

how do we feel about 'making protocols auditable' as an alternative?

this theme includes the issues of how a protocol may be audited against more than one policy, what audit records are required and how they are to be kept, the nature of the criteria for success etc.

bruce

===

From: Prof E Stewart Lee <E.S.Lee@ccsr.cam.ac.uk> Fri Feb 5 1999 15:08
To: B.Christianson@herts.ac.uk, m.roe@ccsr.cam.ac.uk
Subject: Security Protocols Workshop

> this theme includes the issues of how a protocol may be audited against
> more than one policy

This is unwise. Generally, a protocol can enforce more than one policy iff one the policies is a subset of the other. This is a well-known result of composition theory. Policies are even more difficult to compose than objects. Two policies that do not satisfy the subset criterion can sometimes be enforced iff there is a requirement that one be enforced before the other -- e.g.: Mandatory Access Control before Discretionary Access Control. For another example, Bell LaPadula is a policy. So is Non-interference. Their composition is extremely difficult (because NI doesn't have DAC, amongst other more technical reasons).

Stew

X Introductory Remarks

===

From B.Christianson@herts.ac.uk Fri Feb 5 1999 15:06
To: E.S.Lee@ccsr.cam.ac.uk, Michael.Roe@ccsr.cam.ac.uk
Subject: security protocols workshop

i think you are the first to submit a position paper, stew ;-}.

i'm not welded to those particular issues, but the better to inflame
debate:

policy A: (laundry) no payment implies no laundry
policy B: (customer) no laundry implies no payment

neither policy entails the other.

bruce

===

From Michael.Roe@ccsr.cam.ac.uk Fri Feb 5 1999 15:23
To: Prof E Stewart Lee <E.S.Lee@ccsr.cam.ac.uk>
Subject: Re: Security Protocols Workshop

Hmmm ... we're starting to jump into the discussion that should be held
at the workshop here.

It is a fact of life that many protocols involve communication between
parties who have conflicting interests. Each party, if given sole
authority to set policy, would define a policy which conflicts with that
of the other party. Reconciling these conflicts *is* an imporant issue in
protocol design, even though we know that there is no *mathematical* tool
which will cause real conflicts of interest to magically vanish.

I vote we keep Bruce's original sentence.

Mike

Table of Contents

Keynote Address: The Changing Environment	
<i>Roger Needham, Discussion</i>	1
Composing Security Properties	
<i>Stewart Lee, Discussion</i>	6
Auditing against Multiple Policies	
<i>Mark Lomas, Discussion</i>	15
Jikzi: A New Framework for Secure Publishing	
<i>Ross Anderson and Jong-Hyeon Lee</i>	21
<i>Discussion</i>	37
Power and Permission in Security Systems	
<i>Babak Sadighi Firozabadi and Marek Sergot</i>	48
<i>Discussion</i>	54
Auditing against Impossible Abstractions	
<i>Bruce Christianson, Discussion</i>	60
What Is Authentication?	
<i>Dieter Gollmann, Discussion</i>	65
Relations Between Secrets: The Yahalom Protocol	
<i>Lawrence C. Paulson</i>	73
<i>Discussion</i>	78
Modelling Agents' Knowledge Inductively	
<i>Giampaolo Bella</i>	85
<i>Discussion</i>	91
Time-Lock Puzzle with Examinable Evidence of Unlocking Time	
<i>Wenbo Mao</i>	95
<i>Discussion</i>	98
Trust Management and Network Layer Security Protocols	
<i>Matt Blaze, John Ioannidis and Angelos D. Keromytis</i>	103
<i>Discussion</i>	109
Issues in Multicast Security	
<i>Francesco Bergadano, Davide Cavagnino and Bruno Crispo</i>	119
<i>Discussion</i>	132
Performance of Protocols	
<i>Michael Roe</i>	140
<i>Discussion</i>	147

Integrity-Aware PCBC Encryption Schemes
Virgil D. Gligor and Pompiliu Donescu 153
Discussion 169

The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks
Frank Stajano and Ross Anderson 172
Discussion 183

INTERNET-MARKs:
Clear, Secure, and Portable Visual Marks for the Cyber Worlds
Hiroshi Yoshiura, Seiichi Susaki, Yasuhiko Nagai, Tsukasa Saitoh,
Hisashi Toyoshima, Ryoichi Sasaki and Satoru Tezuka 195
Discussion 203

Pictures Can't Lie under Oath
Bruce Christianson, James A. Malcolm and Brian Robinson 208
Discussion 212

Bob versus Bob – Good Guy/Bad Guy
Bob Morris, Discussion 219

Transcript of Concluding Discussion 224

Who Knows?
Allan Ahlberg 228

Author Index 229

Keynote Address: The Changing Environment (Transcript of Discussion)

Roger Needham

Microsoft Research Ltd.

What I wanted to spend a few minutes talking about is changes in the environment and purpose in which protocols are conducted. When Mike Schroder and I wrote our paper in 1977, published in 1978, one of the things we did was to be quite explicit about what we assumed to do with the environment. We assumed that we had principals, whether human or mechanical, that wanted to establish a secure connection, and the assumption was that the principals themselves worked in the way they were supposed to work. If they were people they were honest people, if they were programs they were correct programs. But the environment was thoroughly wicked and everything about the environment conducted itself so as to cause the whole procedure to fail if it possibly could.

That was one assumption. There are a number of other assumptions to do with the environment. For example, that communication was slow, that encryption and decryption was slow, and therefore it was desirable to do them as little as you could; that there were no reliable source of time and therefore that time should not be used whatsoever; and that state was evil. Obviously the principals trying to communicate should set up state, that's what they were trying to do, set up shared state. But servers that help them along should not maintain any state — partly because for reliability, you would have to replicate them and keep them safe and stable in a number of different places, and partly because it would involve things like disc transfers at an embarrassing time, and these were inordinately slow.

Another assumption was that there was no way in which people could have any at all extensive repository of secrets, because the kind of replaceable disc which was released at that time, they were about 2 megabytes capacity and they were about this diameter. They were certainly not something you could put in your pocket or your wallet or your handbag. The other assumption was that computers, although they might be used by one person at a time, they were not personal in the sense that they were owned by an individual. If your machine didn't work, you wandered down the hall until you found somebody who was away, and used the one in that office. They were a serially reusable resource.

Now, it must be obvious that practically every one of the assumptions I have listed isn't true anymore. We do have a reliable source of time; it is possible to maintain personal secrets without it being terribly clumsy; you could if you wanted have a stateful server — non-volatile RAM had not been invented at the time we were writing.

And finally, the assumption that the principals engaging in a communication were persons or things of goodwill and that the rest of the world was evil isn't true either anymore. This is particularly true of electronic commerce, where the

participants in a secure integral transaction may very well have cheating each other as a major goal and not be particularly worried about interference from outside. I would claim that this collection of changes in assumptions adds up to really taking a quite different view from what we do. And I suspect that something that has happened is, that since the study of security of protocols became a cottage industry in the 1980s, people have let the environmental assumptions of the time just become part of the way they think.

It is because the subject has become, in some slightly bad sense, academic. It's practised by a lot of people who actually don't know much about security requirements in the real world, but they can play table-tennis with the symbols and become very skilled at it. So I think it's worth considering — but I have not done this consideration myself — whether we need to revise any of our attitudes.

Another thing that was assumed was that communication was pairwise between an enormous number of principals, one may think of the number of people. This is not to the point either in very many contexts. In many contexts you have organisations which have got an internal network like Microsoft does, which is not particularly secure, it's simply separated. You have communications between people in Microsoft and people in other organisations, say Compaq, where if you wanted to be secure about it, the bit of communication that needs to be secured is the bit that goes from the firewall at the edge of Microsoft to the firewall at the edge of Compaq, and that's going through the wicked world. Now what that has done is to decrease the number of pairs of things that need to have secure connections between them enormously, and it's arguable that this makes possible a completely different approach to doing it. For example, you might say that Microsoft and Compaq equip themselves with a suitably large amount of one-time pad, and every time a visitor physically goes from one place to the other, the visitor carries a few gigabytes of one-time pad in his or her briefcase, and hands it over and it all gets refreshed, and there isn't any infrastructure.

I don't know whether that is particularly sensible, but it's feasible, it's the sort of thing that just wouldn't have been feasible before, and it's worth at any rate thinking about. Similarly, people go on about public key infrastructure. When I used to run the Computer Lab, I thought about what it would be like for us to do our business with our suppliers by electronic commerce. I found that the Lab had about 1800 suppliers which it added to at the rate of one every couple of weeks. We added a customer at the rate of one every five or six weeks. Now these are small numbers, the Computer Lab is not particularly a small organisation and it would be perfectly feasible to make arrangements between it and any of its business correspondents on a pairwise basis, for the Head of Department to write to the company saying we would like to establish an account with you and we will be bound by any electronic signature which is verified by the following public key, puts a stamp on it and sticks it in post, done, no public key infrastructure needed, and you can in fact bootstrap from that up to all sorts of other transactions.

I have not carried any of these thoughts through but I do think that enough has changed, and enough is changing, for it to be appropriate to take a fresh

look at what we do and why we do it, and the circumstances in which we expect it to work, and I shall leave it at that.

Matt Blaze: I'd like to quibble with two things you said, in agreeing with your premise that we need to think about whether the model that we evaluate protocols against is still valid. The first is the virtual private network model that you mentioned, the idea that Microsoft and Compaq talk to each other, so they protect the edges of their network and simplify the problem in that way. Is it actually true that the Microsoft and Compaq networks are more secure against attacks against Microsoft and Compaq than the Internet at large is? I wonder whether the focus on an external threat, which we traditionally looked at in evaluating network security, might be rather foolish. The internal threat may in fact be much greater.

Reply: I think you are very likely right. If you want a really cynical response, it is that corporations behave as if their intranet is secure, and therefore they would feel comfortable in an environment where it was protecting between one corporation and another. Whether they ought to feel comfortable is an entirely different matter, but in practise they do.

Matt Blaze: Because I think internal security will be with us as a requirement for a while. I'd also like to raise the question about the notion of electronic commerce. For a long time we as security people have been talking about the virtues of converting to a paperless workflow system with digital signatures and so on, we've never actively succeeded, but all of a sudden, at least in the United States, we're now seeing laws. States are rushing and competing with each other to see who can be first to pass laws recognising digital signatures as being equivalent to handwritten signatures. And now I think it's gotten away from us. Certainly digital signatures do not have all the same properties as handwritten signatures, and I think it's a tremendous mistake to design security protocols that simply replace the handwritten operations with digital ones, when they're not exactly analogous. In particular, I can't think of an easy way for somebody to steal the ability for me to write my handwritten signature, but I can think of many easy ways they do that with my digital signature.

Reply: I think that's for sure, and my expectation is that, as electronic commerce begins to happen on a big scale, there are going to be some spectacular and ingenious frauds. The present system of doing business with purchase orders and invoices, and copies of the purchase order and cheques and things like that, didn't come into existence overnight, it came into existence over quite a lot of years, with a fair amount of experience of fraud along the way, and I expect the same thing to happen in the electronic world. I think you are absolutely right because people will assume that digital signatures are isomorphic to paper signatures, and that will generate a certain amount of confusion.

Matt Blaze: Some people look forward to it.

Reply: I think some people look forward to it.

Ross Anderson: It may be even worse than that because various things aren't always given the presumption of validity that digital signatures are. This rule erodes the traditional consumer protection that we have in the UK. We have,

for example, credit card legislation similar to your regulation in the States which says that if something goes wrong between about £100 and £3000, then that's the bank's problem not the customer's. If you replace that with a law that says that all digital signatures that are presumed to be valid, then suddenly all the risks are dumped on the poor customer. And so you can expect lots of frauds will follow because of the moral hazard. We think that this is one of the big risks of digital signatures. It's actually coming about because digital signatures have been used as an excuse for a sleight of hand which transfers power from the individuals to the state instead. You don't hear about this at Crypto conferences but perhaps it's an awful lot more important than any of the technical aspects.

Reply: Yes, and it is for reasons connected with that, that when I was saying that something like the Computer Lab might engage in bilateral negotiations with suppliers, I should perhaps have added that I wouldn't expect the Computer Lab to use the same public key for all its suppliers, it might easily have one each. That in itself would be an anti-fraud mechanism, because if somebody steals the signing key they might put themselves in a position to buy a large quantity of toilet paper at the Computer Lab's expense but not to do anything else.

Larry Paulson: I made this remark last year so I'm sorry but I know that in the real world there are signature stamps that are used to sign cheques and there must be controls on those.

Bruce Christianson: Usually they're locked in the safe, but the tremendous advantage is that if somebody's stolen it you can tell — you look in the safe and it's gone.

Stewart Lee: To reinforce what Matt said, somebody — I've forgotten who, it's just fallen out of my mind, it's what happens when you get old — this consultant in New York did a survey of incidents of fraud that involved computers, for security incidents, it involved computers, and discovered that 68% of them were insiders, Bob Courtney, that's his name, and I see no reason at all to expect that the Internet will make any difference in the proportion, about two thirds of them will done by insiders. To reinforce another thing that Matt said, I'm on the Matt bandwagon today, you mentioned that you foresaw that the parts of traffic on the Internet that have to remain secure, could be decreased dramatically if it's intranets talking to intranets sort of thing, but I feel that the Internet is already at a place where there is an enormous number of credit card numbers and expiry dates going over it, and I am very loathe to do that myself because I know how easy it is to observe what goes on, and it is very difficult in fact as a consumer to recover from that. If I order something from some company in Texas and it gets sent to me as information by e-mail, for fifty quid or something like that, it's very difficult to get my fifty quid back, if it's actually somebody else's. How do I prove it?

Reply: Yes, indeed, and what I suspect are the patterns of doing business may change.

David Wheeler: I would speculate that if two thirds of computer frauds are insiders, then two thirds of any fraud are insiders, though it's very easily hushed up.

Stewart Lee: The Bob Courtney survey showed that the small thefts were always publicised, it was the big ones that were hushed up.

Virgil Gligor: Actually, one of the things that Bob showed is that it was more than theft and computer crime. Accidents and mishaps, acts of God, floods, fire, air conditioners blowing up, accounted for most of the losses. Fraud altogether, out of this whole thing, was only about 20% to 23%.

David Wheeler: That's *discovered* fraud.

Virgil Gligor: One thing that bothers me, one difficulty with the Internet. Maybe I'm the only one who's bothered by this. We are locatable at fixed places, we're fixed entities, we are given some sort of Internet address, perhaps two or three, but those are fixed. So we are sitting ducks for people who are willing to attack us. In some sense this happens in reality, we have one place of residence, but generally most of us don't hold much of our valuables in our homes anymore, so we spread our funds in multiple areas, do all sorts of protective things that we learn how to do, and in the Internet we aren't quite so able to do so. In other words, if we were to spread out our wealth of files and data across the Internet somehow, I would feel much better protected. Which reminds me of the paper that Yves Deswarte wrote some years ago in Oakland about fragmenting files and sending them all over the place. Well here is the Internet, we should be able to do that, and we should be able to do a lot more, yet we are not doing it. Does that make any sense to you, is this valuable, is it just a flight of fancy?

Ross Anderson: I think that the Internet threats are wildly overstated by people who are trying to sell the security software. Until a few years ago banks, VISA and so on, were trying to tell everybody: don't do credit card transactions on the net, wait for SET. Now they've got some past history on their books they've changed their mind, and thankfully for them people are beginning to recover from the initial attack of panic. Now I only monitored what was going on closely until last September or so when I handed over the job of reading the relevant banking technology articles to somebody else in the context of Security Reviews, but until then there had been no case anywhere of somebody taking a credit card number off the net. If there had been a case it would have been reported, it would have been all over the headlines. But there was one case where somebody hacked into a merchant server and got credit card numbers that the merchant should not under his merchant agreement ever take. Actually hacking stuff off the net is hard work. We all assume in theory that it can be done but it's not in practice how you would go about getting credit card numbers.

Composing Security Properties

(Transcript of Discussion)

Stewart Lee

Centre for Communications Systems Research
University of Cambridge

What I'm talking about is simultaneously enforcing more than one security policy. What I'm really talking about is when it's over a network where you have two machines that are connected, and you have a security policy in each end. I thought it would be useful to try and decide what we mean by a security policy. It's a directive on the goals for access and integrity and other things, it's an objective or a goal. What I'm going to do in this talk is to very quickly review a number of known security policies. I picked on ones that have been published, that are reasonably well known. I am then going to demonstrate that these policies are inconsistent and any attempt to enforce both of them at the same time is doomed.

The Orange Book is one of the obvious ones. We all like to shoot at it, but there it is, it's a policy and it was embodied in a tract that was first issued in 1983. I very consciously call it a tract because some sort of religious fervour is associated with it. It was followed by a number of other criteria: the British produced the green book, the Canadians produced the blue book, this is brown, the common criteria. That is, the original document that describes it was brown and one must be consistent if nothing else. The Orange Book policy has six fundamental aspects to it: continuous employment of protection; objects have access control labels; individual subjects must be identified. The orange book governs the access to objects by subjects. Objects are passive subjects of actions. People have to be responsible for their actions; it's practical to evaluate the system, in other words it isn't a spaghetti junction of hand written code, and continuous protection must be involved from concept right through to installation.

The Orange Book is modelled on something called the Bell-LaPadula model: reading data at a lower level is allowed, writing data to a higher level is prohibited, destructive writing. You can always write up if it's sending information up but you can't destroy something up there.

Another policy, this time an integrity policy, is a very old paper by Biba. The subject's integrity security level is affected by the integrity security level of the object that it observes. If a subject observes a low integrity object it becomes a low integrity subject, and so on.

Biba has fallen into disuse but it's a policy. Clark-Wilson is a policy that is quite often used. Ross Anderson has used a version of it in medical data. You have unconstrained data items and constrained data items. This is a transaction processing policy for the commercial world and it includes the N-man rule; an event journal, which is where auditing comes in; an integrity verification procedure which was shrouded in mystery in the original paper and still is shrouded

in mystery, and so on. Notice here the terminology is different than in the previous two policies. Not only is the terminology different but this is a policy on a transaction processor that runs on some computer system that has an underlying operating system. Both the previous two have really talked about operating systems.

And finally information flow policies. Information flow policies are governed by some sort of process where there's things with high security level and low security level coming in and things of a high security level and low security level going out. There are two of them that are commonly referred to, one of them is so-called generalised non-interference and it says that the low out cannot be interfered with by either the high in or the high out so that you cannot tell from the low out whether there was any high in or high out there or not. It's not interfered with.

The other one is subtly different. It says nondeducibility and you have exactly the same arrangement but now an observer examining low out cannot deduce with certainty, anything about high in or high out. The key thing here is with certainty, cannot be deduced with certainty. So that something is nondeducibility secure if there is one chance in ten to the something that the message doesn't say what in effect it does say. Non-deducibility security is frequently used in circumstances that involve cryptography because if the process that I spoke about is cryptography then clearly the high in affects the low out.

There are problems when there are several policies and I've got five issues there: information flow versus the Orange Book; Clark-Wilson versus the Orange Book; the Orange Book versus the Orange Book and so on. All of these things have problems of varying difficulty. Information flow really knows only about read up and write down, it doesn't know about what the Orange Book calls discretionary access control at all. The Orange Book has about five other policy things that are just not contained in information flow, so that if one side of the conversation is running information flow and the other is running Orange Book, it won't pass.

Clark-Wilson has the same problem. You have this terminology of constrained data items and unconstrained data items and it doesn't know about a security level. So you have terminology problems. Also, Orange Book doesn't know about N-man rule or UDIs or an integrity verification procedure, and so on.

This is the classic thing where one system has top secret and secret in it, the other system has secret and confidential in it, both are approved for those circumstances, neither of them are approved for the combination of three pieces of data and that is a pretty standard problem.

Orange Book, when doing subject to subject, must be interpreted, because Orange Book really applies to subjects accessing objects. You have to do an interpretation of it, and every time you do an interpretation you're dealing with your own local version of some policy, and that local version may well be inconsistent with others. For instance, when you're dealing with subjects accessing subjects you have flow of control which is an issue. How do you deal with parameters? Is that reading or writing? How do you deal with results? Is that reading or

writing? What do you do about exception handling? These are all places where some things can get into trouble. Sometimes there are other problems. Very few policies other than Clark-Wilson are covered, and it means that one system may need enforcement of something about which the other is ignorant, and that's just not going to work.

Now, something that has to do with inconsistencies. Just to bring us back to earth here, just UNIX against UNIX. It's really not a policy but a mechanism. There are many brands of UNIX and not all of them define the access permissions quite the same way. For instance, the concept of ownership was a little bit different in System V UNIX as against POSIX as against BSD. I presume it's different again in the more modern versions of UNIX. Consequently you have to be very careful of what you do. So what we're trying to do is called composition, we're trying to compose out of several components a system where we can predict the security policy that the system will enforce by knowing the security policies or properties that the components will enforce. There's no known solution to the general composition problem. There is a solution for the general composition problem where the two policies are information flow policies, that is known and was published about a year ago¹. It's known that information flow policies can be composed and you can predict the resulting security policy of the composed system, from the security policy of its components. But that's the only case and it's probably the only case that is practical and considered because the other cases are not sufficiently well analytically modelled to achieve anything.

So what I'm saying here is that two policies cannot be simultaneously enforced unless they are essentially the same policy; or one is a subset of the other and you can use the stronger rather than the weaker; or it's OK to enforce neither of them well; or one of them not well. And this is the point that I wanted to make, and it pertains to this business of auditing the process of a protocol. If you're just auditing for the sake of auditing, keeping a journal of events for the sake of keeping a journal of events, and then auditing that journal of events is busy work involving a great deal of magnetic storage and so on unless you have some objective for doing it. There are a few objectives for doing it, like pinning your finger on who initiated this nonsense, which is reasonable to do, but trying to use this to ensure that some security policies of the end users are being enforced, is unlikely, in the general case at least, to succeed.

I was all inspired by Bruce Christianson. I apologise for the review of well known stuff that was contained in there, but there it was.

Virgil Gligor: Did I understand correctly that you suggest that we use even non composable policies but we audit activities?

Reply: I didn't suggest it, Bruce suggested it.

Bruce Christianson: Stewart initially made an unguarded assertion, which was that you can't enforce two policies simultaneously unless one is a subset of the other. I said, well suppose I have a policy that says you don't get the goods

¹ A. Zakinthinos and E. Stewart Lee, "A Generalised Theory of Security Properties", 1997 Symposium on Security and Privacy, Oakland, California